

Mark My Words!

Linguistic Style Accommodation in Social Media

Cristian Danescu-Niculescu-Mizil*
Cornell University
Ithaca, NY 14853, USA
cristian@cs.cornell.edu

Michael Gamon
Microsoft Research
Redmond, WA 98053, USA
mgamon@microsoft.com

Susan Dumais
Microsoft Research
Redmond, WA 98053, USA
sdumais@microsoft.com

ABSTRACT

The psycholinguistic theory of communication accommodation accounts for the general observation that participants in conversations tend to converge to one another’s communicative behavior: they coordinate in a variety of dimensions including choice of words, syntax, utterance length, pitch and gestures. In its almost forty years of existence, this theory has been empirically supported exclusively through small-scale or controlled laboratory studies. Here we address this phenomenon in the context of Twitter conversations. Undoubtedly, this setting is unlike any other in which accommodation was observed and, thus, challenging to the theory. Its novelty comes not only from its size, but also from the non real-time nature of conversations, from the 140 character length restriction, from the wide variety of social relation types, and from a design that was initially not geared towards conversation at all. Given such constraints, it is not clear a priori whether accommodation is robust enough to occur given the constraints of this new environment. To investigate this, we develop a probabilistic framework that can model accommodation and measure its effects. We apply it to a large Twitter conversational dataset specifically developed for this task. This is the first time the hypothesis of linguistic style accommodation has been examined (and verified) in a large scale, real world setting.

Furthermore, when investigating concepts such as stylistic influence and symmetry of accommodation, we discover a complexity of the phenomenon which was never observed before. We also explore the potential relation between stylistic influence and network features commonly associated with social status.

Categories and Subject Descriptors: J.4 [Computer Applications]: Social and behavioral sciences

General Terms: Measurement, Experimentation, Theory

Keywords: linguistic style accommodation, linguistic convergence, social media, Twitter conversations

*The research described herein was conducted while the first author was a summer intern at Microsoft Research.

Language is a social art.

—Willard van Orman Quine, *Word and Object*

1. INTRODUCTION

The theory of communication accommodation was developed to account for the general observation that in conversations people tend to nonconsciously converge to one another’s communicative behavior: they coordinate in a variety of dimensions including choice of words, syntax, pausing frequency, pitch and gestures [12]. In the last forty years, this phenomenon has received significant attention and numerous studies indicate that such convergence occurs almost instantly for a very diverse set of communication patterns (see Table 1 for examples). These findings suggest that the communicative behavior of conversational partners are “patterned and coordinated, like a dance” [32]. However, up to now this “dance” was exclusively studied in controlled laboratory experiments or through small scale studies. The work presented here demonstrates for the first time the robustness of accommodation theory in a large scale, real world environment: Twitter.

Conversations on Twitter: a new hope. Even though not originally developed as a conversation medium, Twitter turns out to be a fertile ground for dyadic interactions. It is estimated that a quarter of all its users hold conversations with other users on this platform [22] and that around 37% of all tweets are conversational [36]. The fact that these conversations are public renders Twitter one of the largest publicly available resources of naturally occurring conversations.

Undoubtedly, Twitter conversations are unlike those used in previous studies of accommodation. One of the main differences is that these conversations are not face-to-face and do not happen in real-time. Like with email, a user does not need to immediately reply to another user’s message; this might affect the incentive to use accommodation as a way to increase communication efficiency. Another difference is the (famous) restriction of 140 characters per message, which might constrain the freedom one user has to accommodate the other. It is not a priori clear whether accommodation is robust enough to occur under these new constraints.

Also, with very few exceptions, accommodation was only tested in the initial phase of the development of relations between people (i.e., during the acquaintance process) [11]. The relations between Twitter users, on the other hand, are expected to cover a much wider spectrum of development, ranging from newly-introduced to old friends (or enemies).

Thus, also from this perspective, the Twitter environment constitutes a new challenge to the theory.

Linguistic style accommodation. One of the dimensions on which people were shown to accommodate is linguistic style [32, 39, 14], where style denotes the components of language that are unrelated to content: **how** things are said as opposed to **what** is said. This work will focus on this type of accommodation. This is a rather important dimension, since, even though only 0.05% of the English vocabulary is composed of style words (such as articles and prepositions), an estimated 55% of all words people employ are style words [38]. These numbers do not necessarily hold on Twitter, where one might expect style to be sacrificed in favor of content given the length constraint; however, some recent studies also advocate for the importance of style in Twitter [8, 24, 35]. Linguistic style has also been central to a series of NLP applications like authorship attribution and forensic linguistics [30, 41, 18, 23], gender detection [25, 31, 17] and personality type detection [1].

Linguistic style is also known to be, for the most part, generated and processed nonconsciously [26], and thus a suitable vehicle for studying the phenomenon of accommodation, which itself is assumed to occur nonconsciously.

Probabilistic framework. Previous work on accommodation relied mainly on simple correlation-based measures. A new framework is necessary in order to correctly model and measure the effects of linguistic style accommodation in a real world, uncontrolled environment. The main desirables from such a framework are:

- **Comparability:** the effects of accommodation on different components of style should be comparable.
- **Expressivity:** the framework should be expressive enough to permit the evaluation of particular properties of accommodation (discussed in Section 2).
- **Purity:** accommodation should not be confounded with other phenomena.

The last of these desirables is probably the hardest to achieve and thus deserves some discussion here. The main challenge is to distinguish accommodation from the effects of homophily: people that converse are likely to employ a similar linguistic style simply because they know each other. As detailed in Section 5, we control for this effect by using the temporal aspect specific to accommodation: a person can accommodate to her conversational partner only after receiving her input. Another type of potential confusion is that between linguistic style accommodation and topic accommodation; this is avoided in this work by a careful selection of the stylistic features following a methodology employed in psycholinguistic literature (discussed in Section 4).

Stylistic influence and symmetry. Another advantage of the proposed framework is that a new concept of *stylistic influence* emerges naturally: given two conversational partners, one can influence the style of the other more than vice-versa. This concept is a finer-grained version of the concept of *symmetry* of accommodation proposed in the psycholinguistic literature [12]: accommodation can occur symmetrically when both participants in a conversation accommodate

to each other or asymmetrically when only one accommodates. In the latter case, the non-accommodating participant can either maintain her default behavior, or adjust her behavior in the opposite direction from that of the accommodating participant (i.e., *diverge*). We are able to show that imbalance in stylistic influence between Twitter users is preponderant and that symmetry in accommodation is dependent on the stylistic dimension (Figure 4); for example, users are more likely to accommodate symmetrically on the use of 1st person singular pronouns but to accommodate asymmetrically on the use of prepositions. This is the first time such a rich complexity of the accommodation phenomenon is revealed.

A variety of studies relate accommodation and social status. For example, it was hypothesized that a person of lower status will try to accommodate to a person of higher status in order to gain her approval [11, 37]. We take the first steps towards understanding the relation between the concepts of stylistic influence and social status, as reflected in Twitter network features, like number of followers and number of friends, that could be considered (rough) proxies for social status (Section 6.3). Rather surprisingly, we observe almost no correlation between these features and stylistic influence.

Applicability. Apart from its appealing theoretical importance, accommodation also has a variety of potential practical uses. Based on the premise that accommodation has a subtle positive effect on interpersonal communication, Giles et al. [13] discusses applications of accommodation in mediating police-civilian interactions. On a similar note, Taylor and Thomas [39] shows its relevance in the context of hostage negotiations. Accommodation was also shown to be practical in the treatment of mental disability [16] and psychotherapy [9]. In Section 8 we also venture into proposing three new potential applications specific to linguistic style accommodation. We believe that by providing a way to model accommodation and by demonstrating its robustness in a real world environment, the present work provides a framework which supports a wider implementation of such applications.

2. COMMUNICATION ACCOMMODATION THEORY

The psycholinguistic theory of communication accommodation was developed around the following main hypothesis: in dyadic conversations the participants converge to one another’s communicative behavior in terms of a wide range of dimensions [12], both verbal and non-verbal. Table 1 provides a sample of such converging dimensions. Many studies seem to indicate that the communicative behaviors of the participants “are patterned and coordinated, like a dance” [32].

Among various properties of accommodation discussed in the literature, here we briefly review a few that are relevant to our work. First, one should keep in mind that the coordination occurs nonconsciously. Second, accommodation does not necessarily occur *simultaneously* on all dimensions, as shown in [9]. Moreover, convergence on some dimensions does not exclude divergence on others: for example, [3] showed that when conversing with males, females converged on frequency of pauses but diverged on laughter. Another property that is relevant to this work is that of *symme-*

Table 1: Examples of dimensions for which accommodation was observed and the respective studies.

Dimension	Canonical study
Posture	Condon and Ogston, 1967
Pause length	Jaffe and Feldstein, 1970
Utterance length	Matarazzo and Wiens, 1973
Self-disclosure	Derlenga et al., 1973
Head nodding	Hale and Burgoon, 1984
Backchannels	White, 1989
Linguistic style	Niederhoffer and Pennebaker, 2002

try of accommodation: accommodation can occur symmetrically when both participants in a conversation accommodate to each other or asymmetrically when only one accommodates. For example, White [40] presents a study in which Americans accommodate to Japanese on the frequency of backchannels (e.g., ‘hmm’, ‘uh-huh’) but the Japanese did not reciprocate. Asymmetric accommodation has two flavors, depending on the behavior of the non-accommodating participant:

- Default asymmetry: the non-accommodating participant maintains her default behavior (like in the previous example);
- Divergent asymmetry: the non-accommodating participant adjusts her behavior in the opposite direction from that of the accommodating participant (i.e., *diverges*) [12].

It is also worth pointing out that the subject of this work is *instant* accommodation, occurring from one conversational turn to another. Long-term accommodation is considered to be a separate phenomenon with potentially different properties [9, 12]. With a few notable exceptions [9, 32], empirical support for long-term accommodation is absent mostly due to the necessity of longitudinal data.

Various potential explanations for why accommodation occurs have been proposed. One hypothesis is that accommodation occurs from a desire to increase communicational efficiency [37]. Another hypothesis is that a person’s convergence to another person’s communicative patterns is (non-consciously) driven by the desire gain the other’s social approval [11, 37]. Yet another possible motivation is that accommodation is used to “maintain a positive social identity” [20] with the other. The last two hypotheses and several other studies draw a clear relation between social status and accommodation (see also [12]), which will become relevant later in our discussion.

In the present work the focus is on *linguistic style accommodation*, and therefore the work of Niederhoffer and Pennebaker [32] is particularly relevant, being the first study to quantify this phenomenon. It consists of two controlled laboratory experiments (involving 94 dyads) and one study based on transcripts of the Watergate tapes (conversations between Nixon and 3 of his aides) in which coordination on various linguistic style dimensions, like usage of prepositions, adverbs and tentative words is shown to occur between the participants.

In its almost forty years of existence, communication accommodation theory was empirically supported exclusively through small scale studies or controlled laboratory experi-

ments. Also the respective studies focused mainly on real-time interactions (mostly face-to-face, but sometimes computer mediated like in [32]). With this work we aim to change this state of affairs and demonstrate the robustness of this theory in a large scale, real world environment where conversations are not as richly supported as they are in real-time interactions.

3. CONVERSATIONAL DATA

As discussed in Section 1, Twitter is a good environment for our study not only because of its fertility in dyadic interactions, but also because it poses new challenges to the theory of communication accommodation in terms of robustness.

Drawing from this resource, Ritter et al. [36] builds the largest conversational corpus available to date, made up of 1.3 million conversations between 300,000 users. We will refer to this corpus as *conversational dataset A*. In spite of its size, this corpus presents some major drawbacks with respect to the purpose of this paper. First, it has a low density of conversations per pair of conversing users: on average only 4.3 conversations per user; this is not sufficient to model the linguistic style of each pair individually (as required by the accommodation framework proposed in this work and detailed in Section 5). Also, more than half of the pairs of users in this dataset only have unidirectional interaction, i.e., one of the users in a pair never writes to the other. This would not introduce a bias with respect to the type of conversations and relations studied (unidirectional interaction are generally not classified as normal conversations), but would also drastically limit the potential to compare accommodation between users.

To overcome these limitations, we construct a new conversational dataset with very high density of conversations per pair and with reciprocated interactions. We start from conversational dataset A and select all pairs in which both users initiated a conversation at least 2 times. We then collect all tweets posted by these users using the Twitter API¹ and then reconstruct all the conversations between the selected pair. The resulting dataset contains 15 million tweets which make up the complete² public twitter activity (a.k.a. public timeline) of 7,800 users; for each user Twitter metadata (such as the number of friends, the number of followers, the location, etc.) is also available. From these tweets we reconstructed 215,000 conversations between the 2,200 pairs of users with reciprocal relations selected from conversational dataset A, using the same methodology for reconstructing conversations employed in [36]³. This conversational dataset is *complete*, in the sense that all twitter conversations ever held within each pair are available. To the best of our knowledge, this is the largest complete conversational dataset.

The diversity of the user relations and conversations contained in this conversational dataset, dubbed *conversational dataset B.*, is illustrated in the following table summarizing per-pair statistics:

¹<http://apiwiki.twitter.com/>

²Complete up to a maximum 3200 most recent tweets per user, a limitation imposed by the Twitter API.

³Additionally, we remove self replies and retweets from the data on the belief that they do not make part of a proper dyadic interaction.

	Mean	Median	Min	Max
Number of conversations	98	60	1	1744
Average number of turns	2.7	2.6	2	16.8
Days of contact	270	257	1	886

The main unit of interaction in this work is a *conversational turn*, which is defined as two consecutive tweets in a conversation. The two tweets in a turn are always sent by different users and are not re-tweets. Conversational dataset A contains 2.6 million turns and conversational dataset B contains 420,000 turns.⁴

4. MEASURING LINGUISTIC STYLE

Miller [28] shows that style and topic are processed differently in the brain. The distinction between the two is important in our investigation of linguistic style accommodation. In order to measure style and avoid confusion with topic we follow a psycholinguistic methodology used in a variety of applications, known as the LIWC Linguistic Inquiry Word Count (LIWC) method.

LIWC [34] measures word use in psychologically meaningful categories (e.g., articles, auxiliary verbs, positive emotions). It uses over 60 such categories, and dictionaries of words related to each category. This method has been used in a variety of applications (summarized in [38]) including to identify social relations, mental health, and individual traits such as gender, age and relative status. More importantly LIWC is the basis of all recent work on linguistic style accommodation [32, 39, 14] to which we want to relate.

Following the example of these studies, we eliminate all categories related to topic, such as *Leisure*, *Religion* or *Death*. We refer to the 50 remaining dimensions as *style dimensions*. In order to facilitate the presentation of the empirical results, we will focus our discussion on a subset of 14 dimensions that we call *strictly non-topical style dimensions*:

Dimension	Examples	Size
Article	an, the	3
Certainty	always, never	83
Conjunction	but, whereas	28
Discrepancy	should, would	76
Exclusive	without, exclude	17
Inclusive	with, include	18
Indefinite pronoun	it, those	46
Negation	not, never	57
Preposition	to, with	60
Quantifier	few, much	89
Tentative	maybe, perhaps	155
1st person singular pronoun	I, me	12
1st person plural pronoun	we, us	12
2nd person pronoun	you, your	20

For completeness, we mention that all the results presented in this paper also holds for all the other style dimensions (see [38] for a complete list), unless otherwise noted.

We say that a tweet *exhibits* a given stylistic dimension if it contains at least one word from the respective LIWC vocabulary. A tweet can exhibit multiple dimensions and, in fact, the vast majority do.

Although we experimented with different methods of extending the LIWC vocabularies with Twitter-specific expressions, we prefer to keep in line with previous literature on linguistic style matching by using the original vocabularies.

⁴We are unable to make the data public at the time of publication in consideration of the Twitter terms of service.

5. PROBABILISTIC FRAMEWORK

This section introduces a probabilistic framework that can model the phenomenon of accommodation. In defining such a framework, the desirable properties discussed in Section 1 are accounted for: comparability, expressivity and purity. Although designed to be applicable to any type of conversational data and style dimensions, for notational consistency with the rest of the paper, we use the term “tweet” to refer to a conversational utterance.

5.1 Stylistic cohesion

We start by addressing the more general phenomenon of *stylistic cohesion*. It reflects the intuition that tweets belonging to the same conversation are closer stylistically than tweets that do not. Cohesion is defined by comparing the probability that a stylistic dimension is exhibited in tweets that are part of a conversation with the probability that the same dimension is exhibited in unrelated tweets. If the former equals the latter, it means that the distribution of the stylistic dimension is the same whether tweets are part of a conversation or not. If the former is larger than the latter, it means that tweets in a conversation tend to “agree” with respect to the stylistic dimension. If the former is smaller than the latter, it means that tweets in a conversation tend to “disagree” with respect to the stylistic dimension. Formally, for a given dimension C , the measure of stylistic cohesion can be expressed through the following probabilistic expression:

$$Coh(C) \triangleq P(T^C \wedge R^C | T \leftrightarrow R) - P(T^C \wedge R^C) \quad (1)$$

where T^C (respectively R^C) is the event in which a tweet T (respectively R) exhibits C , and $T \leftrightarrow R$ is the condition that tweets T and R form a conversational turn⁵. Thus, demonstrating that cohesion is observable for stylistic dimension C is reduced to showing that $Coh(C) > 0$.

It should be emphasized that accommodation is only one of the possible causes for stylistic cohesion. Another explanation can be the indirect effect of homophily already discussed in Section 1: people that converse are likely to employ a similar linguistic style simply because they know each other or are like each other (we will refer to this as *background* style similarity). This observation motivates the need for a measure which can exclusively target accommodation, discussed next.

5.2 Stylistic accommodation

When defining a probabilistic framework for linguistic style accommodation it is important to control for the effects of *background* style similarity (and provide the purity desiderata introduced in Section 1). Here this is achieved by measuring accommodation for each user pair separately and by taking into account the distinctive temporal nature of accommodation: a user can accommodate to her conversational partner only after receiving her input. In doing so, the concern is eliminated because a confusion with background style similarity effects, like homophily, would not be expected to cause differences within a single pair depending on whether one or the other user in a pair initiates a conversational turn. Therefore, the goal is to measure for a given

⁵The sample space considered throughout this work is the set of all possible ordered conversational tweet pairs.

pair of users a and b who engage in a conversation whether the use of a stylistic dimension C in the initial tweet (of user a) increases the probability of that stylistic dimension in the reply (of user b) beyond what is normally expected from user b (when replying to user a).

Formally, for a given stylistic dimension C and pair of users (a, b) , the accommodation of user b to user a is measured by how much the fact that user a exhibits C in a tweet T_a increases the probability of b to also exhibit C in a reply to T_a :

$$Acc_{(a,b)}(C) \triangleq P(T_b^C | T_a^C, T_b \hookrightarrow T_a) - P(T_b^C | T_b \hookrightarrow T_a) \quad (2)$$

where T_a^C (respectively T_b^C) is the event in which a tweet posted by user a (respectively b) exhibits C , and $T_b \hookrightarrow T_a$ is the condition that T_b is a reply to T_a . This condition, present in both the minuend and the subtrahend⁶, has the role of restricting this measure of accommodation only to replies of b to a , therefore controlling for differences in the background linguistic similarity between users. Also note that by using the \hookrightarrow condition instead of the \leftrightarrow condition employed in the definition of cohesion (1), we embed the distinctive temporal aspect of accommodation mentioned earlier. This ability to integrate temporal disparity is an essential advantage of this framework over the correlation based measures previously used in studies of stylistic accommodation [32, 39, 14]⁷.

Since the main goal is to address global accommodation (as opposed to the within-pair accommodation described above), the accommodation for a given dimension C is defined as:

$$Acc(C) = E[Acc_{(a,b)}(C)] \quad (3)$$

where the expectation is taken over all possible conversing pairs (a, b) . Under this framework, proving that accommodation is observable for stylistic dimension C is reduced to showing that $Acc(C) > 0$.

5.3 Stylistic influence and symmetry

One important property of the way Acc is defined is its asymmetry: the accommodation of user b to user a on stylistic dimension C is potentially different from the accommodation of user a to user b on the same stylistic dimension. The notion of stylistic influence arises naturally:

$$I_{(a,b)}(C) \triangleq Acc_{(a,b)}(C) - Acc_{(b,a)}(C) \quad (4)$$

for a given stylistic dimension C . If $I_{(a,b)}(C) > 0$ we can say that b accommodates more to a on C than b does to a .

A related concept is accommodation symmetry (discussed in Section 2), which is tied to the accommodation measure in the following way. Given that b accommodates to a , i.e. $Acc_{(a,b)}(C) > 0$, we have

- Symmetry when $Acc_{(b,a)}(C) > 0$,
- Default asymmetry when $Acc_{(b,a)}(C) = 0$,
- Divergent asymmetry when: $Acc_{(b,a)}(C) < 0$

⁶Where by minuend we mean the left term of a subtraction and by subtrahend we mean to the right one.

⁷As a concrete example, correlation does not distinguish between the case in which the initial tweet exhibits C but the reply does not, and the reverse case in which the initial tweet does not exhibit C but the reply does.

6. EMPIRICAL VALIDATION

Equipped with the probabilistic framework introduced in the previous section, here we proceed with an empirical validation of the accommodation phenomenon on the conversation data described in Section 3. As previously discussed, this setting is fundamentally different from all other circumstances in which the theory of communication accommodation was validated, therefore challenging its robustness.

6.1 Validation of stylistic cohesion

We start by asking whether Twitter conversations are characterized by stylistic cohesion, since this is a precondition for accommodation. The stylistic cohesion model described in Section 5.1 does not distinguish between users and therefore can be directly applied to the conversational dataset A (introduced in Section 3).

In order to demonstrate that cohesion is exhibited in our data we estimate the two probabilities involved in (1) as follows. We estimate the first probability as the fraction of all turns in which both tweets exhibit dimension C :

$$\hat{P}(T^C \wedge R^C | T \leftrightarrow R) = \frac{|\{(t, r) | t \leftrightarrow r, t^C, r^C\}|}{|\{(t, r) | t \leftrightarrow r\}|} \quad (5)$$

where t^C denotes the condition that a tweet t exhibits C .⁸

To estimate the second probability, we first construct a set of “fake turns” by randomly pairing together tweets from the entire conversational data (regardless of their authors). We can then write:

$$\hat{P}(T^C \wedge R^C) = \frac{|\{(t, r) | t \not\leftrightarrow r, t^C, r^C\}|}{|\{(t, r) | t \not\leftrightarrow r\}|} \quad (6)$$

where t^C is the condition that the tweet t exhibits C and $t \not\leftrightarrow r$ is the condition that the tweets t and r are paired together in a fake turn.

Establishing that cohesion is exhibited in the data corresponds to rejecting the null hypothesis of these two probabilities being equal. Fisher’s exact test⁹ rejects this hypothesis with p-value smaller than 0.0001 for each of the strictly non-topical style dimensions. Figure 1 shows the estimates of the two probabilities for each of these style dimensions (the difference between the two is shown in red/dark). While this result is not surprising, it is a necessary precondition for verifying the more subtle hypothesis of accommodation that we are going to address next.

6.2 Validation of stylistic accommodation

We now proceed to answer the main question of this work: does the hypothesis of stylistic accommodation proposed in the psycholinguistic literature hold in social media conversations? Since the probabilistic framework for accommodation is applied at the level of user pairs, the conversational dataset B is employed for this analysis.

For each ordered user pair (a, b) and stylistic dimension C , we estimate the minuend in (2) as the fraction of b ’s replies to a in which b ’s tweet t_b exhibits C :

⁸Lowercase letters are used to represent tweets that make up our dataset, distinguish them from the uppercase letters that refer to probabilistic events in the framework defined in Section 5.

⁹We use this exact variant of the χ^2 test since for some style dimensions the expected counts are low.

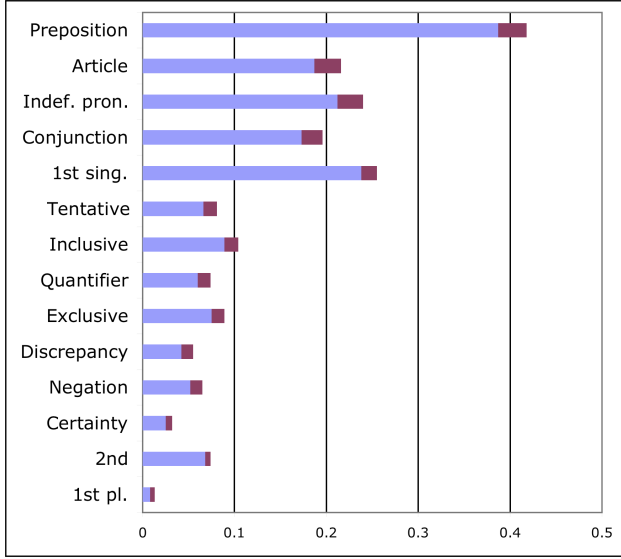


Figure 1: The effect of stylistic cohesion observed as the difference between $\hat{P}(T^C \wedge R^C | T \leftrightarrow R)$ (composite bars) and $\hat{P}(T^C \wedge R^C)$ (blue bars). The differences, shown in red/dark, are statistically significant ($p < 0.0001$). The dimensions are shown in decreasing order of the difference.

$$\hat{P}(T_b^C | T_b \leftrightarrow T_a) = \frac{|\{(t_a, t_b) | t_b \leftrightarrow t_a, t_b^C\}|}{|\{(t_a, t_b) | t_b \leftrightarrow t_a\}|} \quad (7)$$

Similarly, the subtrahend is estimated as:

$$\hat{P}(T_b^C | T_a^C, T_b \leftrightarrow T_a) = \frac{|\{(t_a, t_b) | t_b \leftrightarrow t_a, t_b^C, t_a^C\}|}{|\{(t_a, t_b) | t_b \leftrightarrow t_a, t_a^C\}|} \quad (8)$$

We can then measure the amount of accommodation $\widehat{Acc}(C)$ exhibited in our dataset as the difference between the mean of the set of subtrahend estimations

$$\left\{ \hat{P}(T_b^C | T_a^C, T_b \leftrightarrow T_a) \mid (a, b) \in \text{Pairs} \right\}$$

and the mean of the minuend estimations

$$\left\{ \hat{P}(T_b^C | T_b \leftrightarrow T_a) \mid (a, b) \in \text{Pairs} \right\},$$

where Pairs is the set of all ordered pairs¹⁰. Figure 2 compares these means — the former is illustrated in red/right, the latter in blue/left — for each strictly non-topical stylistic dimension. All the differences are statistically significant with a p-value smaller than 0.0001 according to a two-tailed paired t-test¹¹ for all strictly non-topical style dimensions with the exception of the *2nd person pronoun* stylistic dimension for which the difference is not statistically significant.

Even though our focus is on the strictly non-topical style dimensions, for completeness we also measured accommoda-

¹⁰We discard all user pairs for which the denominator of any of these two estimations is less than 10.

¹¹In order to allay concerns regarding the independence assumption of this test, for each two users a and b we only consider one of the two possible ordered pairs (a, b) and (b, a) .

tion on the remaining 36 dimensions and observed a statistically significant effect for all of them except for *Fillers* (like ‘blah’, ‘yaknow’) for which the data was insufficient.

Note that by design, our probabilistic framework allows comparison between the accommodation effects exhibited for each dimension C (i.e., fulfills the comparability desiderata introduced in Section 1). Here are some of the comparisons worth pointing out:

- Users accommodate significantly more on *tentativeness* than on *certainty* (p-value smaller than 0.01 according to an independent t-test).¹²
- Users accommodate significantly more on *negative emotions* than on *positive emotions* (not illustrated, $\widehat{Acc}(\text{Neg. emo.}) = 0.07$, $\widehat{Acc}(\text{Pos. emo.}) = 0.04$; p-value smaller than 0.01 according to an independent t-test for the difference).
- *1st person singular pronoun* vs. *2nd person pronoun*. In retrospect, the fact that accommodation is not exhibited for the *2nd person pronoun* dimension seems natural: words like ‘you’ have a different meaning for two participants involved in a conversation. However, the same holds for the *1st person singular pronoun* dimension for which accommodation is observed. This could be explained by the social-psychology hypothesis of disclosure reciprocity in dyadic relationships [6].

With the results presented here we are able to verify that accommodation does indeed hold in large scale, real world conversational setting with properties that a priori seemed challenging to the theory. In the remainder of this section we will use our framework to investigate what properties linguistic style accommodation exhibits in this conversational setting.

6.3 Stylistic influence and symmetry

Here we seek to understand the role that the concept of stylistic influence (introduced in Section 5.3) has in Twitter conversations. We start by asking whether stylistic influence is prevalent in the data: in general, is there a balance between the amount two participants in a conversation accommodate? Or, on the contrary, is one user stylistically dominating the other?

In terms of our framework, we can test whether in expectation there is an imbalance of accommodation between participants in a conversation by verifying whether we can reject the null hypothesis $E[\text{abs}(I_{(a,b)}(C))] = 0$, where the expectation is taken over all conversing pairs (a, b) . Using definition (4), this is reduced to rejecting:

$$E[\text{abs}(Acc_{(a,b)}(C) - Acc_{(b,a)}(C))] = 0.$$

and further to rejecting:

$$E[\max(Acc_{(a,b)}(C), Acc_{(b,a)}(C))] = E[\min(Acc_{(a,b)}(C), Acc_{(b,a)}(C))]$$

where the first term is the expected accommodation of the most accommodating users (where the accommodation is

¹²Therefore doubt appears to be more “contagious” than confidence.




Figure 2: The effect of accommodation $\widehat{Acc}(C)$ for each strictly non-topical stylistic dimension C observed as the difference between the means of $\{\widehat{P}(T_b^C | T_b, T_a) \mid (a, b) \in \text{Pairs}\}$ (blue, left) and $\{\widehat{P}(T_b^C | T_a^C, T_b, T_a) \mid (a, b) \in \text{Pairs}\}$ (red, right). All the differences are statistically significant ($p < 0.0001$), except for the *2nd person pronoun* category. The dimensions are ordered according to the amount of accommodation observed.

always compared within each pair), and can be estimated the mean of:

$$\left\{ \max \left(\widehat{Acc}_{(a,b)}(C), \widehat{Acc}_{(b,a)}(C) \right) \mid (a, b) \in \text{Pairs} \right\},$$

and the second term is the expected accommodation of the least accommodating users, estimated by the mean of:

$$\left\{ \min \left(\widehat{Acc}_{(a,b)}(C), \widehat{Acc}_{(b,a)}(C) \right) \mid (a, b) \in \text{Pairs} \right\}.$$

Using the same method for estimating $Acc_{(a,b)}(C)$ discussed in Section 6.2, we reject this hypothesis for all strictly non-topical style dimensions C (paired t-test with p-value smaller than 0.0001)¹³. Figure 3 illustrates the difference between the expected accommodation of the least accommodating users (red/left) and that of the most accommodating users (blue/right) in a pair. A difference in the type of imbalance between dimension is revealed; for example, while for *1st person plural pronouns* in general the least accommodating users still match the style of the most accommodating participants (even though significantly less than vice-versa), for *certainty* the least accommodating users in general diverge from the style of the most accommodating participants.

To further investigate this intriguing difference between style dimensions, we turn our attention to the property of symmetry. Figure 4 shows the percentage split between symmetrically accommodating pairs (blue/left), asymmetrically default accommodating pairs (yellow/center) and asymmetrically diverging accommodating pairs (red/right), as defined in Section 5.3.

The conclusions that can be drawn from analyzing these results is that accommodation is a much more complex behavior than previously reported in the literature, where it was assumed that only one type of accommodation occurs

for a given dimension¹⁴. But as it can be observed in Figure 4 all three types of accommodation have a considerable stake. Furthermore, in all previous work on linguistic style accommodation, no distinction was made between the type of accommodation occurring for each dimension. However, our study indicates a clear difference between dimensions:

- Symmetric accommodation is dominant for *1st pron. pl.*, *Discrepancy* and *Indef. pron.*;
- Asymmetric accommodation (of both types) is dominant in most of the other dimensions;
- Asymmetric diverging accommodation is dominant for *2nd person pronoun*.

A potential explanation for the fact that such a complex accommodation behavior was not previously observed may be the difference between the Twitter conversational setting and that traditionally used in the literature (discussed in Section 1), especially in the spectrum of relation types covered (mostly limited to one type in the previous studies). Another explanation may be the increased expressibility of our probabilistic framework over the correlation based framework used in previous studies.

6.4 Relation to social status

As pointed out in Section 2, the psycholinguistic literature draws clear a connection between the social status of a user and its tendency to accommodate. Therefore, it is natural to ask whether stylistic influence correlates with differences in social status between the users and we take the first steps to address this question. For lack of a better proxy, we employ user features that were previously reported to be related to social influence on Twitter [2]. For each pair of users

¹³The same holds for all the other dimensions except *Fillers* for which the data was insufficient.

¹⁴Here we refer to any dimension of accommodation, like the ones in Table 1, not only to linguistic style dimensions.




Figure 3: The effect of stylistic influence for each strictly non-topical stylistic dimension C observed as the difference between the means of $\left\{ \min \left(\widehat{Acc}_{(a,b)}(C), \widehat{Acc}_{(b,a)}(C) \right) \mid (a,b) \in \text{Pairs} \right\}$ (red, left) and $\left\{ \max \left(\widehat{Acc}_{(a,b)}(C) \mid (a,b) \in \text{Pairs} \right) \right\}$ (blue, right). All the differences are statistically significant ($p < 0.0001$). The dimensions are shown in decreasing order of the difference.

in our data we compare: #followers, #followees, #posts, #days on Twitter, #posts per day and ownership of a personal website. We find that for all style dimensions none of these features correlate strongly with stylistic influence; the largest positive Pearson correlation coefficient obtained was 0.15 between #followees and stylistic influence on *1st pron. pl.*. Also, for the task of predicting the most influential user in each pair a decision tree classifier¹⁵ rendered relatively poor results. The best improvement over the majority class baseline was of only 7% for the *1st pron. pl.* dimension (in this case the most predictive features were the difference in #friends and the difference in ownership of a personal website). All this suggests that stylistic influence appears to be only weakly connected to these social features. However, one should take this observation with a grain of salt: the proxies for social status available on Twitter and employed here are far from ideal. Future work should seek to use better proxies for social status, possibly in environments with richer social data.

7. RELATED WORK

Here we briefly touch on related work not already discussed. Much of the research in understanding social media focuses on the network relations between users. More recently, this line of work has been complimented with a rich analysis of the content of posts as well as structural relations among posters. In one early study combing these two dimensions of analysis, Paolillo [33] examined linguistic variations associated with strong and weak ties in an early internet chat relay system. The strength of friendship ties on Facebook was related by Gilbert and Karahalios [10] to various language features including intimacy words, positive and negative emotions. Eisenstein et al. [8] investigated the role geographic variation of language has in Twitter and Kiciman [24] examined the extent to which differences in language models of Twitter posts (as measured by perplexity) were related to metadata associated with the senders.

¹⁵We used the Weka implementation of the C4.5 decision tree, available at www.cs.waikato.ac.nz/ml/weka/




Figure 4: The percentage of accommodating pairs that exhibit each of the three types of accommodation: symmetric, default asymmetry and diverging asymmetry.

Such work demonstrates the importance of linguistic style variations in Twitter which also plays a crucial role in our study.

Latent variable models have also been used to summarize more general linguistic patterns in social media. Ramage et al. [35] developed a partially supervised learning model (Labeled LDA) to summarize key linguistic trends in a large collection of millions of Twitter posts. They identified four general types of dimensions, which they characterized as substance, status, social and style. These included dimensions about events, ideas, things, or people (substance), related to some socially communicative end (social), related to per-

sonal updates (status), and indicative of broader trends of language use (style). This representation was used to improve filtering of tweets and recommendations of people to follow. In the task of tweet ranking, a different approach is taken by [7] which employs twitter specific features in conjunction to textual content. Another way of characterizing key trends in text data is to use known distinctions or dimensions. In addition to the already discussed work based on LIWC, see [29, 33] for other examples of analyses of linguistic variation with respect to position or status in a social network. Of particular interest in our work is the distinction between linguistic style and content. Style or function words make up about 55% of the words that we speak, read and hear according to Tausczik and Pennebaker [38], similar to findings of Ramage et al. [35] in their analysis of Twitter. In our research, we use LIWC to characterize the linguistic style of posts as well as individuals.

One particularly interesting type of linguistic activity in social media has to do with conversations, that is with exchanges between one or more individuals. Twitter conversations are the main focus in this work. Java et al. [22] found that 21% of users in their study used Twitter for conversational purposes (as measured by the use of @, a convention to address a post to a particular user), and that 12.5% of all posts were part of conversations. Honeycutt and Herring [19] analyzed conversational exchanges on the Twitter public timeline, focusing on the function of the @ sign. They found that short dyadic conversations occur frequently, along with some longer multi-participant conversations. Ritter et al. [36] developed an unsupervised learning approach to identify conversational structure from open-topic conversations. Specifically they trained an LDA model which combined conversational (speech acts) and content topics on a corpus of 1.3 million Twitter conversations, and discovered interpretable speech acts (reference broadcast, status, question, reaction, comment, etc.) by clustering utterances with similar conversational roles. In our research, we build on this data set and extend it to include the complete conversational history of individuals over a period of almost one year.

Since the notion of linguistic style is central to this work, we also want to point out other instances in which it plays an important role. Linguistic style was shown to be crucial in the area of authorship attribution and in forensic linguistics (for an overview see [23]). To identify an author, it is necessary to look beyond content into the — often subconscious — stylistic properties of the language. Simple measures like word length, word complexity, sentence length and vocabulary complexity were at the forefront of earlier research into attribution problems (e.g. [41, 18]). Since Mosteller and Wallace’s seminal work on the Federalist Papers [30], however, a trend has emerged to focus on the distribution of function words as a diagnostic for authorship, a method that in various incarnations now dominates the research. Other areas using similar methods include gender detection from text [25, 31, 17] and personality type detection [1].

8. CONCLUSIONS AND FUTURE WORK

In this paper we have shown that the hypothesis of linguistic style accommodation can be confirmed in a real life, large scale dataset of Twitter conversations. We have developed a probabilistic framework that allows us to measure accommodation and, importantly, to distinguish effects of style accommodation from those of homophily and

topic-accommodation. We also have demonstrated how this framework allows us to formalize and investigate the notions of stylistic influence and asymmetric accommodation.

It is important to highlight that our findings are anything but obvious, given that Twitter is a medium unlike any other setting in which the phenomenon was previously observed. Its novelty comes not only from its size, but also from the wide variety of social relation types, from the non real-time nature of conversations, from the 140 character length restriction and from a design that was initially not geared towards conversation at all. This work demonstrates that accommodation is robust enough to occur under these new constraints, presumably because it is deeply ingrained in human social behavior.

We believe that this line of research has a number of natural extensions. One question we have not addressed is the issue of long-term accommodation: can we measure accommodation over a longer period of time, from the first interaction of two users on? Answering this question is challenging because it requires richer longitudinal data. It would also be very interesting to explore interplay between the accommodating behavior and the type of social relation.

As for practical applications, on the premise that accommodation renders conversations more pleasant and effective, we posit that having the linguistic style of automated dialogue systems match that of the user would increase the quality of the interaction. Personalized ranking of tweets could also benefit by selecting tweets with styles that match that of the tweets issued by the target user. Finally, given the evidence that this work brings to support the universality of the accommodation phenomenon, we envision its use in detection of forged conversations.¹⁶

Finally, we hope that our findings will stimulate further research and refinements of the communication accommodation theory in the psycholinguistic world.

Acknowledgments We thank Lillian Lee for inspiring conversations, Munmun De Choudhury, Scott Counts, Sumit Basu, Dan Liebling, Magdalena Narożniak, Alexandru Niculescu-Mizil, Tim Paek, Bo Pang, Chris Quirk for technical advice and the anonymous reviewers for helpful comments. This paper is based upon work supported in part by the NSF grant IIS-0910664.

9. REFERENCES

- [1] S. Argamon, S. Dhawle, and M. Koppel. Lexical predictors of personality type. *Joint Annual Meeting of the Interface and the Classification Society of North America*, 2005.
- [2] E. Bakshy, J. M. Hofman, W. A. Mason, and D. J. Watts. Everyone’s an influencer: Quantifying influence on Twitter. *WSDM*, 2011.
- [3] F. Bilous and R. Krauss. Dominance and accommodation in the conversational behavior of

¹⁶We are inspired here by the use of Benford’s law in detecting forged financial reports [4]. Though potentially not as common as such forgery, situations in which conversational transcripts are contested are not infrequent. One recent example is the October 2010 release of phone conversations between top Romanian political leaders and a compromised business man. Another one is the controversy surrounding the reality TV shows “The Jersey Shore” and “The Hill” which are suspected of being scripted.

- same- and mixed-gender dyads. *Language and Communication*, 8, 183-194, 1988.
- [4] W. T. Cho and B. Gaines. Breaking the (Benford) law. *The American Statistician*, 61(3):218-223, 2007.
- [5] Condon and Ogston. A segmentation of behavior. *Journal of psychiatric research*, 1973.
- [6] V. Derlega, M. Harris, and A. Chaikin. Self-disclosure reciprocity, liking and the deviant. *Journal of Experimental Social Psychology*, 9(4):277-284, 1973.
- [7] Duan, Jiang, Qin, Zhou, and Shum. An empirical study on learning to rank of tweets. *COLING*, 2010.
- [8] J. Eisenstein, B. O'Connor, N. A. Smith, and E. P. Xing. A latent variable model for geographic lexical variation. *EMNLP*, 2010.
- [9] K. Ferrara. Accommodation in therapy. In *Contexts of accommodation: developments in applied sociolinguistics*. Cambridge University Press, 1991.
- [10] E. Gilbert and K. Karahalios. Predicting tie strength with social media. *HCI*, 2009.
- [11] H. Giles. Communication accommodation theory. In *Engaging theories in interpersonal communication: multiple perspectives*. Sage Publications, 2008.
- [12] H. Giles, J. Coupland, and N. Coupland. Accommodation theory: Communication, context, and consequences. In *Contexts of accommodation: developments in applied sociolinguistics*. Cambridge University Press, 1991.
- [13] H. Giles, M. Willemyns, C. Gallois, and M. Anderson. Accommodating a new frontier: The context of law enforcement. *K. Fiedler (Ed.), Social Communication. New York: Psychology Press.*, 2006.
- [14] A. L. Gonzales, J. T. Hancock, and J. W. Pennebaker. Language style matching as a predictor of social dynamics in small groups. *Communication Research*, 37(1):3-19, Feb 2010.
- [15] J. Hale and J. Burgoon. Models of reactions to changes in nonverbal immediacy. *Journal of Nonverbal Behavior*, 8(4):287-314, 1984.
- [16] H. Hamilton. Accommodation and mental disability. In *Contexts of accommodation: developments in applied sociolinguistics*. Cambridge University Press, 1991.
- [17] S. Herring and J. Paolillo. Gender and genre variation in weblogs. *Journal of Sociolinguistics*, Jan 2006.
- [18] D. Holmes. Authorship attribution. *Computers and the Humanities*, 28(2):87-106, Apr 1994.
- [19] C. Honeycutt and S. Herring. Beyond microblogging: Conversation and collaboration via twitter. *HICSS*, 2009.
- [20] D. A. Infante, A. S. Rancer, and D. F. Womack. *Building communication theory*. Waveland Press, 2006.
- [21] J. Jaffé and S. Feldstein. *Rhythms of dialogue*. Academic Press, 1970.
- [22] A. Java, X. Song, T. Finin, and B. Tseng. Why we twitter: understanding microblogging usage and communities. *WebKDD/SNA-KDD*, Aug 2007.
- [23] P. Juola. *Authorship Attribution*. Now Publishers, 2008.
- [24] E. Kiciman. Language differences and metadata features on Twitter. *Web N-gram Workshop*, 2010.
- [25] M. Koppel, S. Argamon, and A. Shimoni. Automatically categorizing written texts by author gender. *Literary and Linguistic Computing*, 2002.
- [26] W. Levelt and S. Kelter. Surface form and memory in question answering. *Cognitive Psychology*, 14(1):78-106, 1982.
- [27] J. Matarazzo and A. Wiens. *Interview: Research on Its Anatomy and Structure*. Chicago: Aldine, 1973.
- [28] G. Miller. *The science of words*. Scientific American Library. Scientific American Library, 1996.
- [29] L. Milroy and J. Milroy. Social network and social class: Toward an integrated sociolinguistic model. *Language in Society*, 21(1):1-26, Mar 1992.
- [30] F. Mosteller and D. Wallace. Inference in an authorship problem. *Journal of the American Statistical Association*, 58(302):275-309, Jun 1963.
- [31] A. Mukherjee and B. Liu. Improving gender classification of blog authors. *EMNLP*, 2010.
- [32] K. Niederhoffer and J. Pennebaker. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 2002.
- [33] J. Paolillo. Language variation on internet relay chat: A social network approach. *Journal of Sociolinguistics*, 2001.
- [34] J. W. Pennebaker, R. J. Booth, and M. E. Francis. Linguistic Inquiry and Word Count (LIWC): A computerized text analysis program. *LIWC.net*, 2007.
- [35] D. Ramage, S. Dumais, and D. Liebling. Characterizing microblogs with topic models. *International AAAI Conference on Weblogs and Social Media*, 2010.
- [36] A. Ritter, C. Cherry, and B. Dolan. Unsupervised modeling of Twitter conversations. *NAACL*, 2010.
- [37] R. L. Street and H. Giles. Speech accommodation theory. In *Social cognition and communication*. Sage Publications, 1982.
- [38] Y. R. Tausczik and J. W. Pennebaker. The psychological meaning of words: LIWC and computerized text analysis methods. *Journal of Language and Social Psychology*, 29(1):24-54, Mar 2010.
- [39] P. Taylor and S. Thomas. Linguistic style matching and negotiation outcome. *Negotiation and Conflict Management Research*, 1(3):263-281, 2008.
- [40] S. White. Backchannels across cultures: A study of Americans and Japanese. *Language in Society*, 18(1):59-76, Mar 1989.
- [41] G. Yule. On sentence-length as a statistical characteristic of style in prose: With application to two cases of disputed authorship. *Biometrika*, 30(3/4):363-390, Jan 1939.