

Alignment, Acceptance, and Rejection of Group Identities in Online Political Discourse

Hagyeong Shin

Dept. of Linguistics
San Diego State University
San Diego, CA, USA, 92182
hshin@sdsu.edu

Gabriel Doyle

Dept. of Linguistics
San Diego State University
San Diego, CA, USA, 92182
gdoyle@sdsu.edu

Abstract

Conversation is a joint social process, with participants cooperating to exchange information. This process is helped along through linguistic alignment: participants' adoption of each other's word use. This alignment is robust, appearing many settings, and is nearly always positive. We create an alignment model for examining alignment in Twitter conversations across antagonistic groups. This model finds that some word categories, specifically pronouns used to establish group identity and common ground, are negatively aligned. This negative alignment is observed despite other categories, which are less related to the group dynamics, showing the standard positive alignment. This suggests that alignment is strongly biased toward cooperative alignment, but that different linguistic features can show substantially different behaviors.

1 Introduction

Conversation, whether friendly chit-chat or heated debate, is a jointly negotiated social process, in which interlocutors balance the assertion of one's own identity and ideas against a receptivity to the others. Work in Communication Accommodation Theory has demonstrated that speakers tend to converge their communicative behavior in order to achieve social approval from their in-group members, while they tend to diverge their behavior in a conversation with out-group members, especially when the group dynamics are strained (Giles et al., 1991, 1973).

Linguistic alignment, the use of similar words to one's conversational partner, is one prominent and robust form of this accommodation, and has been detected in a variety of linguistic interactions, ranging from speed dates to the Supreme Court (Danescu-Niculescu-Mizil et al., 2011; Guo et al., 2015; Ireland et al., 2011; Niederhoffer and

Pennebaker, 2002). In particular, this alignment is usually positive, reflecting a widespread willingness to accept and build off of the linguistic structure provided by one's interlocutor; the differences in alignment have generally been of degree, not direction, subtly reflecting group differences in power and interest.

The present work proposes a new model of alignment, SWAM, which adapts the WHAM alignment model (Doyle and Frank, 2016). We examine alignment behaviors in a setting with clear group identities and enmity between the groups but with uncertainty on which group is majority or minority: conversations between supporters of the two major candidates in the 2016 U.S. Presidential election. Unlike previous alignment work, we find some cases of substantial negative alignment, especially on personal pronouns that play a key role in assigning group identity and establishing common ground in the discourse. In addition, within-versus cross-group conversations show divergent patterns of both overall frequency and alignment behaviors on pronouns even when the alignment is positive. These differences contrast with the relatively stable (though still occasionally negative) alignment on word categories that reflect possible rhetorical approaches within the discussions, suggesting that group dynamics within the argument are, in a sense, more contentious than the argument itself.

2 Previous Studies

2.1 Linguistic Alignment

Accommodation in communication happens at many levels, from mimicking a conversation partner's paralinguistic features to choosing which language to use in multilingual societies (Giles et al., 1991). One established approach to assess accommodation in linguistic representation

is to look at the usage of function word categories, such as pronouns, prepositions, and articles (Danescu-Niculescu-Mizil et al., 2011; Niederhoffer and Pennebaker, 2002). This approach argues that function words provide the syntactic structure, which can vary somewhat independently of the content words being used. Speakers can express the same thought through different speech styles and reflect their own personality, identity, and emotions (Chung and Pennebaker, 2007).

In this context, we limit our analysis to convergence in lexical category choices, which can be the consequence of both social and cognitive processes. We call this specific quantification of accommodation “linguistic alignment”, but it is closely related to general concepts such as priming and entrainment. This alignment behavior may be the result of social or cognitive processes, or both, though we focus on the social influences here.

2.2 Linguistic Alignment between Groups

Recent models of linguistic alignment have attempted to separate homophily, an inherent similarity in speakers’ language use, from adaptive alignment in response to a partner’s recent word use (Danescu-Niculescu-Mizil et al., 2011; Doyle et al., 2017). If homophily is not separated from alignment, it is impossible to compare within- and cross-group alignment, since the groups themselves are likely to have different overall word distributions. Both alignment and homophily can be meaningful; Doyle et al. (2017) combine the two to estimate employees’ level of inclusion in the workplace.

Separating these factors opens the door to investigate alignment behaviors even in cases where different groups speak in different ways; if homophily is not factored out, cross-group differences will produce alignment underestimates. Thus far, these models of alignment have been applied mostly in cases where there is a single salient group that speakers wish to join (Doyle et al., 2017), or where group identities are less salient than dyadic social roles or relationships, such as social power (Danescu-Niculescu-Mizil et al., 2012), engagement (Niederhoffer and Pennebaker, 2002), or attraction (Ireland et al., 2011).

There is some evidence and an intuition that alignment can cross group boundaries, but it has not been measured using such models of adaptive linguistic alignment. Niederhoffer and Pen-

nebaker (2002) pointed out that speakers with negative feelings are likely to coordinate their linguistic style to each other, while speakers who are not engaged to each other at all are less likely to align their linguistic style. Speakers also might actively coordinate their speech to their opponents’ in order to persuade them more effectively (Burlison and Fennelly, 1981; Duran and Fusaroli, 2017). If two people with different opinions are talking to each other, they may also align their speech style as a good-faith effort to understand the other’s position (Pickering and Garrod, 2004).

However, it is also reasonable to expect that speakers with enmity would diverge their speech style as a way to express their disagreement to each other, especially if they feel disrespected or slighted (Giles et al., 1991). At the same time, if the function word usage can reflect speakers’ psychological state (Chung and Pennebaker, 2007), then negative alignment to opponents would be observed as a fair representation of the disagreement between speakers. Supporting this idea, Rosenthal and McKeown (2015) showed that accommodation in word usage could be a feature to improve their model detecting agreement and disagreement between speakers.

In the present work, we consider cross-group alignment on personal pronouns, which can express group identity, as well as on word categories that may indicate different rhetorical approaches to the argument (Pennebaker et al., 2003). Van Swol and Carlson (2017) suggests that the pronoun category can be useful markers of group dynamics in a debate setting, and Schwartz et al. (2013) suggests that it is reasonable to expect the different word usage from different groups. In fact, although we find mostly positive alignment, we do see negative alignment in some cross-group uses, suggesting strong group identities can overrule the general desire to align.

3 Data

3.1 Word categories

This study examines alignment and baseline word use on 8 word categories from Linguistic Inquiry and Word Count (LIWC; Pennebaker et al. (2007)), a common categorization method in alignment research. Details on word categories and example words for each category can be found in Table 1. For example, the first person singular pronoun *I* category counted 12 different forms of

Category	Example	Size
1st singular (<i>I</i>)	I, me, mine	12
2nd person (<i>You</i>)	you, your, thou	20
1st plural (<i>We</i>)	we, us, our	12
3rd plural (<i>They</i>)	they, their, they'd	10
Social processes	talk, they, child	455
Cognitive processes	cause, know, ought	730
Positive emotion	love, nice, sweet	406
Negative emotion	hurt, ugly, nasty	499

Table 1: Word categories for linguistic alignment with examples and the number of word tokens in the category.

the *I* pronoun, such as *I, me, mine, myself, I'm, and I'd*.

We choose four pronoun categories (*I, you, we, they*) to investigate the relationship between group dynamics and linguistic alignment. We expect that in a conversation between in-group members, *I, we, they* will be observed often. When these pronouns are initially spoken by a speaker, repliers can express their in-group membership while aligning to their usage of the words at the same time. In the conversation with out-group members, *you* usage will be observed more often because it will allow repliers to refer to the speaker while excluding themselves as a part of the speaker's group. In the cross-group conversation, alignment on inclusive *we* indicates that repliers acknowledged and expressed themselves as a member of speakers' in-group. However, alignment on exclusive *they* in cross-group conversation should be interpreted with much more attention. When a replier is aligning their usage of *they* to their out-group member, it likely indicates that both groups are referring to a shared referent, implying enough cooperation to enter an object into common ground (Clark, 1996).

Additionally, four rhetorical word categories are considered. In LIWC, psychological processes are categorized into social processes, cognitive processes, and affective processes, the last of which covers positive and negative emotions. Social and affective process categories are, as their names indicate, the markers of social behavior and emotions. Cognitive process markers include words that reflect causation (*because, hence*), discrepancy (*should, would*), certainty (*always, never*), and inclusion (*and, with*), to name a few. A speaker's baseline usage of rhetorical categories

will present the group-specific speech styles that may be dependent on group identity, reflecting preferred styles of argument. The degree of alignment on rhetorical categories indicates whether speakers maintain their group's discussion style or adapt to the other group.

3.2 Twitter Conversation

The corpus data was built specifically for this research. The population of the data was Twitter conversations about the 2016 presidential election dated from July 27th, 2016 (a day after both parties announced their candidates) to November 7th, 2016 (a day before the election day). Twitter users were divided into two different groups according to their supporting candidates, based on the assumption that all speakers included in the data were partisans and had a single supporting candidate. When the users' supporting candidate was not explicitly shown in their speech, additional information was considered, including previous Tweets, profile statements, and profile pictures. Speakers' political affiliation was first coded by the researcher and the coder's reliability was tested. Two other coders agreed on the researcher's coding of 50 users (25 were coded as Trump supporters and 25 were coded as Clinton supporters) with Fleiss' Kappa score 0.87 ($\kappa = 0.86, p < 0.001$) with average 94.4% confidence in their answers.

3.3 Sampling Method

The corpus data was built by a snowball method from seed accounts. Seed accounts spanned major media channels (@cnnbrk; @FoxNews; @NBCNews; @ABC) and the candidates' Twitter accounts (@realDonaldTrump; @HillaryClinton). The original Twitter messages from the seed accounts were not considered as a part of the data, but replies and replies to replies were. The minimal unit of the data was a paired conversation extracted from the comment section. An initial message *a* (single Twitter message, known as a tweet) and the following reply *b* created a pair of the conversation.

3.4 Datasets

In total, four sets of Twitter data were gathered. The first two datasets (TT, CC) consisted of conversations between members of the same group (within-group conversation). The other two datasets (TC, CT) consisted of conversations

Dataset	Message	Reply
TT	I saw a poll where she was up by 8 here and people say that they hate her so who knows	I'm not sure what's going on with that, but #Reality tells a different story. Someone's lying.
	Dems spreading lies again. They are projecting Trump as stupid but voters knows he is not!!	Dems running out of ways to get rid of Trump, so now they will push this BS that he is CRAZY and out of control!!
CC	Jill Stein is more concerned about Hillary than she is about Trump. That tells you all you need to know about this loon. #Green-TownHall	You make a lot of sense. I'm sick of Hillary bashing.
TC	#Libtard you are in for the shock of your life #TRUMPTRAIN	You are if you think he has a chance; have you any idea what the Electoral U map looks like?... Hillary!
	Unbelievable that #TeamHillary thinks America is so stupid we won't notice that moderator is a close Clinton pal	Not America. Just folks like you. #Trumpsuneducated
CT	Fight - not a fight - he was told Mexico won't pay for it. Why should they?	Never expected Mex 2 write us a check. Other ways 2 make them pay for wall. Trump knows how 2 negotiate.

Table 2: Examples of conversation pairs from each dataset (First letter indicates initiator's group, second indicates replier's)

across the groups (cross-group conversation). In the dataset references, Trump supporters' message is represented with T, and Clinton supporters' message is represented with C. The first letter indicates the initiator's group; the second indicates the replier's group. There is an average of 266 unique repliers in each group.

4 SWAM Model

This study adapts the Word-Based Hierarchical Alignment Model (WHAM; Doyle and Frank (2016)) to estimate alignment on different word categories in the Twitter conversations. WHAM defines two key quantities: *baseline* word use, the rate at which someone uses a given word category W when it has not been used in the preceding message, and *alignment*, the relative increase in the probability of words from W being used when the preceding message used a word from W .

Both quantities have been argued to be psychologically meaningful, with baseline usage reflecting internalization of in-group identity, homophily, and enculturation, and alignment reflecting a willingness to adjust one's own behavior to fit another's expectations and framing (Doyle et al., 2017; Giles et al., 1991).

The WHAM framework uses a hierarchy of nor-

mal distributions to tie together observations from related messages (e.g., multiple repliers with similar demographics) to improve its robustness when data is sparse or the sociological factors are subtle. This requires the researcher to make statistical assumptions about the structure's effect on alignment behaviors, but can improve signal detection when group dynamics are subtle or group membership is difficult to determine (Doyle et al., 2016).

However, when the group identities are strong and unambiguous, this inference can be excessive, and may even lead to inaccurate estimates, as the more complex optimization process may create a non-convex learning problem. The Bayesian hierarchy in WHAM also aggregates information across groups to improve alignment estimates; in cases where the groups are opposed, one group's behavior may not be predictive of the other's. We propose the Simplified Word-Based Alignment Model (SWAM) for such cases, where group dynamics are expected to provide robust and possibly distinct signals.

WHAM infers two key parameters: η^{align} and η^{base} , the logit-space alignment and baseline values, conditioned on a hierarchy of Gaussian priors.




Figure 1: Generative schematic of the Simplified Word-Based Alignment Model (SWAM) used in this study. Hierarchical parameter chains from the WHAM model are eliminated; alignment values are fit independently by word category and conversation group.

SWAM estimates the two parameters directly as:

$$\eta^{base} = \log p(B|notA) \quad (1)$$

$$\eta^{align} = \log \frac{p(B|A)}{p(B|notA)}, \quad (2)$$

where $p(B|A)$ is the probability of a replier using a word category when the initial message contained it, and $p(B|notA)$ is the probability of the replier using it when the initial message did not.

SWAM treats alignment as a change in the log-odds of a given word in the reply belonging to W , depending on whether W appeared in the preceding message. SWAM can be thought of as a midpoint between WHAM and the subtractive alignment model of Danescu-Niculescu-Mizil et al. (2011), with three main differences from the latter model. First, SWAM’s baseline is $p(B|notA)$, as opposed to unconditioned $p(B)$ for Danescu-Niculescu-Mizil et al. (2011). Second, SWAM places alignment on log-odds rather than probability, avoiding floor effects in alignment for rare word categories. Third, SWAM calculates by-word alignment rather than by-message, controlling for the effect of varying message/reply lengths. These three differences allow SWAM to retain the improved fit of WHAM (Doyle et al., 2016), while gaining the computational simplicity and group-dynamic agnosticism of Danescu-Niculescu-Mizil et al. (2011).

5 Results

5.1 Pronouns

The results of baseline frequency and alignment values for the four conversation types are presented in Figure 2 and 3, respectively. We analyze each pronoun set in turn.

First of all, baseline usage of *you* shows that *you* was used more often among repliers in the cross-group conversations. However, the alignment pattern for *you* was much stronger in within-group conversations. That is, repliers are generally more likely to use *you* in cross-group settings to refer to out-group members overall, but within the group, one member using *you* encourages the other to use it as well.

You alignment in within-group conversation could reflect rapport-building, a sense that speakers understand each other well enough to talk about each other, and an acceptance of the other’s common ground (as in the example for CC in Table 2). On the other hand, *you* alignment in between-group conversations should be interpreted as the result of disagreement to each other (See examples for TC in Table 2). *You* alignment in this case is the action of pointing fingers at each other, which happens at an overall elevated level, regardless of whether the other person has already done so.

Baseline usage of *they* shows the opposite pattern from *you* usage, with higher *they* usage in the in-group conversations. This type of *they* usage can be a reference to out-group members (see the second example for TT in Table 2). By using *they*, repliers can express their membership as a part of the in-group and make assertions about the out-group. It also can reflect acceptance of the interlocutor placing objects in common ground, which can be referred to by pronouns.

They alignment patterns were comparable across the conversation types, except that Trump supporters showed divergence when responding to Clinton supporters. The CT conversation in Table 2 reflects this divergence, with *Mexico* being repeated rather than being replaced by *they*, suggesting Trump supporters reject the elements Clinton supporters attempt to put into common ground.

Moving on to baseline usage of *we*, Trump supporters were most likely to use this pronoun, especially in their in-group conversations, suggesting a strong awareness of and desire for group identity. Contrary to the alignment patterns of *they*, Clin-




Figure 2: Repliers' baseline usage of category markers is the probability of usage of the word when it has not been said by the initial speaker.

ton supporters were actively diverging their usage of *we* from Trump supporters. Meanwhile, Trump supporters were not actively diverging on *we* as they did for the *they* usage.

Claiming in-group membership by using in-group identity marker can be one way of claiming common ground, which indicates that speakers belong to the group who shares specific goals and values (Brown and Levinson, 1987). Therefore, Trump supporters' baseline use and alignment of *we* and *they* suggest that they were accepting and reinforcing common ground with in-group members by using *we*, but rejecting common ground with out-group members by not aligning to *they*. Clinton supporters showed a different way of reflecting their acceptance and rejection. They chose to reject common ground by not aligning to their out-group members' in-group marker *we*, but seemed to accept the common ground within the conversation built by out-group members' use of *they*.

Interestingly, *I* showed the least variability, both in baseline and alignment, across the groups. However, *I* is also the only one of these pronoun groups that does not refer to someone else, and thus should be least affected by group dynamics. In fact, we see Chung and Pennebaker (2007)'s general finding of solid *I*-alignment, even in cross-group communication.




Figure 3: Repliers' alignment on category markers represents the probability of repliers' usage of the word when it has been said by the initial speaker.

Overall, we see effects both in the baseline and alignment values that are consistent with a strong group-identity construction process. Furthermore, we see strong negative alignment in cross-group communication on pronouns tied to group identity and grounding, showing that cross-group animosity can overrule the general pattern of positive alignment in certain dimensions. However, the overall alignment is still positive; even the rejection of certain aspects of the conversation do not lead to across-the-board divergence.

5.2 Rhetorical Categories

Despite our hypothesis that the rhetorical categories of words could indicate different groups' preferred style of argumentation, these categories showed limited variation compared to the pronouns. The baseline values only varied a small amount between groups, with Clinton supporters having slightly elevated baseline use of social and cognitive words, and slightly less positive emotion.

The alignment values were mostly small positive values, much as has been observed in stylistic alignment in previous work. However, cross-group Trump-Clinton conversations did have negative alignment on cognitive processes. This category spans markers of certainty, discrepancy, and inclusion, and has been argued to reflect argumentation framing that appeals to rationality. This may

be a sign of rejecting or dismissing their interlocutors' argument framing. But overall, there is no strong evidence of differences in alignment in argumentative style in this data, and the bulk of the effect remains on group identification.

A possible reason for the lack of differences in argumentation style may be uncertainty about the setting of the cross-group communication. Elevated causation word usage has been argued to be employed by the minority position within a debate, to provide convincing evidence against the status quo (Pennebaker et al., 2003; Van Swol and Carlson, 2017). The datasets consist of conversations from the middle of the election campaign, when it was uncertain which group was in the majority or minority (as seen in the first TT conversation in Table 2). This uncertainty may have led both groups to adopt more similar argumentation styles than if they believed themselves to occupy different points in the power continuum.

6 Discussion

From our results, we see that social context affected pronoun use and alignment, which fits into the Communication Accommodation Theory account (Giles et al., 1991). Meanwhile, rhetorical word use and alignment was independent of social context between speakers, though it is unclear whether this reflects a perception of equal footing in their power dynamics or is driven primarily by automatic alignment influences rather than social factors (Pickering and Garrod, 2004). To expand the scope of this argument, we can further test if the negative alignment can be found in other LIWC categories as well, which have no clear group-dynamic predictions.

One thing to point out is that even though pronouns and some rhetorical words are categorized as function words, which have been hypothesized to reflect structural rather than semantic alignment (Chung and Pennebaker, 2007), these category words are still somewhat context- and content-oriented. That is, use and alignment of some function words is inevitable for speakers to stay within the topic of conversation or to mention the entity whose referential term is already set in the common ground. From Trump supporters' negative alignment on *they*, we could see that speakers were in fact able to actively reject the reference method by not using the content-oriented function words. In the future work, it will be meaningful to sepa-

rate the alignment motivated by active acceptance and agreement from the alignment that must have occurred in order to stay within the conversation.

Testing our hypotheses in different settings can help to resolve this issue. One possibility is to separate the election debate into small sets of conversations with different topics, and then compare the alignment patterns between sets. Because of the lexical coherence that each topic of conversations have, we will be able to better separate the effect of context- and content-oriented words from the linguistic alignment result. As a result, we might be able to see negative alignment on rhetorical category between subset of conversations. We can also test our hypotheses with different languages. Investigating alignment in languages that do not use pronouns heavily for reference can be useful to see how the group dynamics are expressed through different word categories. Particles in some languages, such as Japanese and Korean, can mark specific argument roles, and this linguistic structure can allow us to detect syntactic alignment without looking much into the context- and content-oriented function words. Lastly, the SWAM model is an adaptation of the WHAM model, and while the basic patterns look similar to those found by WHAM, a more precise comparison of the models' estimates with a larger dataset is an important step to ensure that the SWAM estimates are accurate.

7 Conclusion

Pronoun usage and alignment reflect the group dynamics between Trump supporters and Clinton supporters, and observations of negative alignment are consistent with a battle over who defines the groups and common ground. However, the use and alignment of rhetorical words were not substantially affected by the group dynamics but rather reflected that there was an uncertainty about who belongs to the majority or minority group. In a political debate or conversation between opponents, speakers are likely to project their group identity with the usage of pronouns but are likely to maintain their rhetorical style as a way to maintain their group identity.

References

Penelope Brown and Stephen C. Levinson. 1987. *Politeness: Some universals in language usage*, volume 4. Cambridge university press.

- Brant R. Burleson and Deborah A. Fennelly. 1981. The effects of persuasive appeal form and cognitive complexity on children's sharing behavior. *Child Study Journal*, 11:75–90.
- Cindy Chung and James W Pennebaker. 2007. The psychological functions of function words. *Social communication*, 1:343–359.
- Herbert H. Clark. 1996. *Using Language*. Cambridge University Press, Cambridge.
- Cristian Danescu-Niculescu-Mizil, 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*, pages 745–754.
- Cristian Danescu-Niculescu-Mizil, 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*, pages 699–708.
- Gabriel Doyle and Michael C. Frank. 2016. Investigating the sources of linguistic alignment in conversation. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 526–536.
- Gabriel Doyle, Amir Goldberg, Sameer Srivastava, and Michael Frank. 2017. Alignment at work: Using language to distinguish the internalization and self-regulation components of cultural fit in organizations. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 603–612.
- Gabriel Doyle, Dan Yurovsky, and Michael C. Frank. 2016. A robust framework for estimating linguistic alignment in Twitter conversations. In *Proceedings of the 25th international conference on World Wide Web*, pages 637–648.
- Nicholas D. Duran and Riccardo Fusaroli. 2017. Conversing with a devils advocate: Interpersonal coordination in deception and disagreement. *PloS one*, 12(6):e0178140.
- Howard Giles, Nikolas Coupland, and Justine Coupland. 1991. Accommodation theory: Communication, context, and consequences. In Howard Giles, Justine Coupland, and Nikolas Coupland, editors, *Contexts of accommodation: Developments in applied sociolinguistics*, pages 1–68. Cambridge University Press, Cambridge.
- Howard Giles, Donald M. Taylor, and Richard Bourhis. 1973. Towards a theory of interpersonal accommodation through language: Some canadian data. *Language in society*, 2(2):177–192.
- Fangjian Guo, Charles Blundell, Hanna Wallach, and Katherine Heller. 2015. The Bayesian Echo Chamber: Modeling Social Influence via Linguistic Accommodation. In *Proceedings of the 18th International Conference on Artificial Intelligence and Statistics*, pages 315–323.
- Molly E. Ireland, Richard B. Slatcher, Paul W. Eastwick, Lauren E. Scissors, Eli J. Finkel, and James W. Pennebaker. 2011. Language style matching predicts relationship initiation and stability. *Psychological Science*, 22:39–44.
- Kate G. Niederhoffer and James W. Pennebaker. 2002. Linguistic style matching in social interaction. *Journal of Language and Social Psychology*, 21(4):337–360.
- James W. Pennebaker, Roger J. Booth, and Martha E. Francis. 2007. *Linguistic Inquiry and Word Count: LIWC*. Austin, TX: *liwc.net*.
- James W. Pennebaker, Matthias R. Mehl, and Kate G. Niederhoffer. 2003. Psychological aspects of natural language use: Our words, our selves. *Annual review of psychology*, 54(1):547–577.
- Martin J. Pickering and Simon Garrod. 2004. Toward a mechanistic psychology of dialogue. *Behavioral and brain sciences*, 27(2):169–190.
- Sara Rosenthal and Kathy McKeown. 2015. I couldnt agree more: The role of conversational structure in agreement and disagreement detection in online discussions. In *Proceedings of the 16th Annual Meeting of the Special Interest Group on Discourse and Dialogue*, pages 168–177.
- H. Andrew Schwartz, Johannes C. Eichstaedt, Margaret L. Kern, Lukasz Dziurzynski, Stephanie M. Ramones, Megha Agrawal, Achal Shah, Michal Kosinski, David Stillwell, Martin E.P. Seligman, and Lyle H. Ungar. 2013. Personality, gender, and age in the language of social media: The open-vocabulary approach. *PloS one*, 8(9):e73791.
- Lyn M. Van Swol and Cassandra L. Carlson. 2017. Language use and influence among minority, majority, and homogeneous group members. *Communication Research*, 44(4):512–529.