

A COMPUTATIONAL APPROACH TO LINGUISTIC COORDINATION

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Cristian Danescu-Niculescu-Mizil

August 2012

© 2012 Cristian Danescu-Niculescu-Mizil
ALL RIGHTS RESERVED

A COMPUTATIONAL APPROACH TO
LINGUISTIC COORDINATION

Cristian Danescu-Niculescu-Mizil, Ph.D.

Cornell University 2012

Many online human activities leave digital traces that are recorded in natural-language format. The exploitation of this unprecedented resource under a computational framework can bring a phase transition in our understanding of human social behavior and shape the future of social media systems. This thesis describes a computational approach to an intriguing aspect of conversational behavior, *linguistic style coordination*: conversational participants tend to immediately and unconsciously adapt to each other's language styles; a speaker will even adjust the number of articles and other function words in their next utterance in response to the number in their partner's immediately preceding utterance.

We develop a probabilistic framework that can model and quantify this phenomenon and apply it to four different conversational settings: Twitter conversations, discussions between Wikipedia editors, oral arguments in the U.S. Supreme Court and movie script dialogs. We are able to examine linguistic style coordination in large scale settings for the first time, and by doing so, to reveal important properties that were never observed in previous studies, which were conducted mainly in small and controlled environments.

The resulting new understanding of conversational behavior bears practical importance in the task of uncovering key properties of social relations. We present an analysis framework based on linguistic coordination that can be used to shed light on power relationships and that works consistently across multiple types of power — including a

more “static” form of power based on status differences, and a more “situational” form of power in which one individual experiences a type of dependence on another.

Finally, our computational framework can be employed to gain insight into the causal mechanism behind coordination behavior. It was previously hypothesized in the psycho-linguistic literature that coordination has arisen as a way to achieve social goals, such as gaining approval or emphasizing difference in status. But has the adaptation mechanism become so deeply embedded in the language-generation process as to become a reflex? We argue that fictional dialogs offer a way to study this question, since authors create the conversations but don’t receive the social benefits (rather, the imagined characters do). Indeed, we find significant coordination across many families of function words in our large movie-script corpus. We thereby provide evidence that language coordination is so implanted within our conception of conversational behavior that, even if such coordination is socially motivated, it is exhibited even when the person generating the language in question is not receiving any of the presumed social advantages.

BIOGRAPHICAL SKETCH

In second grade, Cristian wrote a short autobiographical note in a letter addressed to the “children of the world”, dated June 4th 1991:

I study all day during my occupied time, while in my free time I watch TV, read, play and sometimes write computer code. [...] I want to become:

1. an athlete; 2. a cosmonaut; 3. an electrician;¹ 4. an inventor. [...] And I wish to study well. I don’t study too well, but not too bad either.

Cristian continued to study — not too well, but not too bad either — up to the day in which he wrote this biographical sketch, for a total of 23 years: eight years of elementary and middle school in “Scoala 19”, four years of high-school in the special-math group of “Colegiul National Mihai Viteazul”, four years of undergraduate studies at the University of Bucharest in Romania, one of which was spent as an Erasmus exchange student at the University of Bari in Italy, two years of Master’s at Jacobs University Bremen in Germany and five years at Cornell University where he completed his Ph.D. degree. After all this, Cristian would still like to become some kind of an inventor and to wear a white lab coat from once in a while.

¹Translator’s note: Cristian was confused with regard to what being an “electrician” meant, perhaps being an “electrical engineer” was closer to what he probably meant to write.

To my brother Alex,
who has taught me how to walk, read, multiply, code and perform cross-validation.

ACKNOWLEDGEMENTS

I was immensely fortunate during my five years at Cornell. No words can describe the gratitude I feel towards Lillian Lee, my advisor and my friend. I really did fear her as a tyrant, love her as a mother, revere her as a god, but I also think she is funny.²

Another one of my life's great fortunes was to work with Jon Kleinberg — although I really want to write “to think with Jon Kleinberg”, to follow his beautiful reasoning, to attempt to add my humble contributions, or to dare to emulate it: “how would Jon think about this?”.

There are many others that contributed to my wonderful experience: Judith Bernstock, my history of art minor advisor; the exceptional people I co-authored with: Andrei Broder, Justin Cheng, Eunsol Choi, Rick Ducott, Susan Dumais, Evgeniy Gabrilovich, Michael Gamon, Vanja Josifovski, Jon Kleinberg, Gueorgi Kossinets, Lillian Lee, Bo Pang, Chenhao Tan and Mark Yatskar; my part-time officemates: Steven An and Jeff Chadwick, who were in the office just the right amount of time; my next-door-officemates, who adored my impromptu visits so much: Art Munson, Myle Ott and Vasu Raman; the old guard, from whom I learned the tricks of the trade: Eric Breck, Yejin Choi, Huijia Lin, Tudor Marian and Filip Radlinski; and Becky Stewart, for helping with basically everything.

Also, there was much outside Upson Hall contributing to me enjoying every day in Ithaca: the history of art classes, the wine class, the Italian/French/German language dinners, the faculty lunch leftovers, the soups at the Temple of Zeus, the salads at the Terrace, the windsurfing club, the Barton Ballers, the Megahurts, the second dam, my jogging trail, the Treman Willow, my porch, my friends.

I am very grateful to my mom and dad for joining me in this endeavor, for believing in me and for giving me lots of last names. I also thank Diana Minculescu for being next

²I am paraphrasing here one of our favorite New Yorker's cartoon, authored by Zachary Kanin.

to me during my last days in Ithaca. And, finally, I thank Catalin Draghici for sharing with me the geometry problem solution he dreamt of in the 7th grade, this thesis would have never existed if I would have not “stolen the solution” from him.

The work presented in this thesis was supported in part by the National Science Foundation grants ISS-0910664 and IIS-1016099, grants from Google and Yahoo! and a Yahoo! Key Scientific Challenges award. Part of this research was conducted while I was an intern at Microsoft Research.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
1 Introduction	1
1.1 The Conversational Coordination Effect	1
1.2 Linguistic Style Coordination	1
1.3 Coordination In Online Settings	3
1.4 Coordination And Power Differences	4
1.5 Causal Mechanism: Social Strategy Or Reflex	4
2 Measuring linguistic coordination	6
2.1 Linguistic Style Markers	6
2.2 Coordination Measure	7
2.3 Properties Of The Coordination Measure	8
2.3.1 Bounds	8
2.3.2 Discussion of alternative measures	9
2.4 Stylistic Influence	10
2.5 Coordination Towards A Group	11
2.6 Aggregated Measures	12
3 Coordination in social media interactions	14
3.1 Twitter Conversational Dataset	15
3.2 Empirical Validation	16
3.2.1 Validation of stylistic coordination	17
3.2.2 Stylistic influence and symmetry	18
3.2.3 Potential alternative explanations	21
4 Coordination and power differences	26
4.1 Linguistic Coordination And Power	28
4.2 Power Relations In Wikipedia And Supreme Court Data	30
4.2.1 Discussions among Wikipedia editors	31
4.2.2 Supreme Court oral arguments	32
4.2.3 Power Relations in the Data	34
4.3 Empirical Investigation	35
4.3.1 Power from status: Wikipedia	36
4.3.2 Power from status: Supreme Court	41
4.3.3 Power from dependence	42
4.4 Cross-Domain Analysis	44
4.5 Interactions Among Hypotheses \mathcal{P} And \mathcal{B}	47

5	Causal Mechanism	51
5.1	Movie Dialogs Corpus	54
5.2	Experimental Results	54
5.2.1	Coordination exists in fictional dialogs	54
5.2.2	Self-coordination	57
5.2.3	Coordination and imagined relation	58
6	Further related work	64
6.1	Linguistic Style And Human Characteristics	64
6.2	Language In Social Media	64
6.3	Power, Structural Features And Language	66
6.4	Collaborative Communities	67
6.5	Imagined Conversations	67
7	Summary and future directions	69

CHAPTER 1

INTRODUCTION

1.1 The Conversational Coordination Effect

The conversational coordination effect is the “nonconscious mimicry of the postures, mannerisms, facial expressions, and other behaviors of one’s interaction partners” (Chartrand and Bargh, 1999).¹ For example, if one conversational participant crosses their arms, their partner often unconsciously crosses their arms as well. The effect occurs for language, too, ranging from matching of acoustic features such as accent, speech rate, and pitch (Chartrand and van Baaren, 2009; Giles et al., 1991) to lexico-syntactic priming across adjacent or nearby utterances (Bock, 1986; Pickering and Garrod, 2004; Reitter et al., 2011; Ward and Litman, 2007). Table 1.1 provides a sample of dimensions of interaction on which people were found to converge, together with the respective representative studies. The emergence of this phenomenon on a wide variety of dimensions suggests that communicative behaviors of the participants become “patterned and coordinated, like a dance” (Niederhoffer and Pennebaker, 2002).

1.2 Linguistic Style Coordination

One of the dimensions on which people were shown to coordinate is linguistic style (Gonzales et al., 2010; Niederhoffer and Pennebaker, 2002; Taylor and Thomas, 2008), where style denotes the components of language that are unrelated to content: **how** things are said as opposed to **what** is said. This is a rather important dimension, since

¹Related terms include adaptation, alignment, entrainment, priming, chameleon effect and Du Bois’ dialogic syntax (Du Bois, 2001).

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

Table 1.1: Examples of dimensions for which coordination was observed and the respective studies.

even though only 0.05% of the English vocabulary are style words (such as articles and prepositions), an estimated 55% of all words people employ are style words (Tausczik and Pennebaker, 2010). Linguistic style has also been central to a series of NLP applications like authorship attribution and forensic linguistics (Holmes, 1994; Juola, 2008; Mosteller and Wallace, 1963; Yule, 1939), gender detection (Herring and Paolillo, 2006; Koppel et al., 2002; Mukherjee and Liu, 2010; Sarawgi et al., 2011) and personality type detection (Argamon et al., 2005; Mairesse and Walker, 2011).

The work presented in this thesis focuses on *linguistic style coordination*, specifically adjacent-utterance coordination with respect to classes of function word. To exemplify the phenomenon, we discuss two short conversations.

- *First example:* The following exchange from the movie “The Getaway” (get, 1972) demonstrates quantifier coordination.

Doc: At least you were outside.

Carol: It doesn’t make much difference where you are [...]

Note that “Carol” used a quantifier, one that is different than the one “Doc” employed. Also, notice that “Carol” could just as well have replied in a way that doesn’t include a quantifier, for example, “It doesn’t really matter where you are...”.

- *Second example:* This example comes from an experiment involving preposition coordination (Levelt and Kelter, 1982). Shopkeepers who were called and asked “At what time does your shop close?” were significantly more likely to say “At five o’clock” than “five o’clock”.²

1.3 Coordination In Online Settings

In the last forty years, this linguistic “dance” was observed and studied almost exclusively in small-scale or controlled laboratory studies. A priori, it is not all clear whether the phenomenon is robust enough to occur under the constraints imposed by online social networks, where most conversations are not face-to-face, do not happen in real-time and are subject to various stylistic constraints. Furthermore, no formalism was proved suitable to study this phenomenon in natural conversational settings. We address this issue in Chapter 2 where we develop a probabilistic framework that can model coordination and measure its effects in a large scale, real-life environment (Danescu-Niculescu-Mizil et al., 2011).

This framework allows us to analyze the coordination phenomenon in the context of Twitter conversations, using a large dataset specifically developed for this task, and to show for the first time that linguistic style coordination emerges in a large scale, uncontrolled setting. Using our probabilistic framework to exploit the richness of the Twitter conversational setting, we were also able to provide an unprecedented understanding of the qualitative aspects of the intriguing phenomenon. For example, we were able to show that there is often a coordination imbalance between participants in a conversation: very often one interlocutor is stylistically dominating the other.

²This is an example of lexical matching manifested as part of syntactic coordination.

1.4 Coordination And Power Differences

But what factors mediate this asymmetry between the levels of coordination? How do characteristics of the users and of their relationships affect their coordination behavior? In Chapter 4 we focus on power relations and show how these are reflected in coordination behavior (Danescu-Niculescu-Mizil et al., 2012b). We find that, in general, people with low power exhibit greater language coordination than people with high power; conversely, people coordinate more with interlocutors who have higher power than with those who have lower power.

The relation between status level and the extent of language coordination holds across multiple types of power — including a more “static” form of power based on status differences, and a more “situational” form of power in which one individual experiences a type of dependence on another — and across two very different domains: the Wikipedia community of editors and the U.S. Supreme Court oral arguments.

1.5 Causal Mechanism: Social Strategy Or Reflex

All the results discussed above bring a new understanding of the properties of coordination behavior, without, however, providing any explanation of the causal mechanism behind this intriguing phenomenon. This is an issue highly debated in the psycholinguistic literature (Branigan et al., 2010; Ireland et al., 2011). One line of thought is that coordination represents a social strategy whose aim is to gain the other’s social approval (Giles, 2008; Street and Giles, 1982) or enhance the other’s comprehension (Bell, 1984; Bortfeld and Brennan, 1997; Clark, 1996).

But an important question is whether the adaptation mechanism has become so

deeply embedded in the language-generation process as to have transformed into a reflex not requiring any social triggering. In Chapter 5 we are able to gain new insights by taking a radical approach to this problem: we consider a setting in which the persons *generating* the coordinating dialog are different from those *engaged* in the dialog (and standing to reap the social benefits) — imagined conversations, specifically, scripted movie dialogs (Danescu-Niculescu-Mizil and Lee, 2011).

CHAPTER 2

MEASURING LINGUISTIC COORDINATION

In this chapter we introduce a new framework for measuring linguistic style coordination by quantifying the degree to which one individual *immediately* echoes the linguistic style of the person they are responding to. Here, linguistic style is quantified by a person’s usage of certain linguistic style markers (corresponding to categories of function words). We start with a description of these markers, followed by a formal definition of coordination. We then discuss properties of the coordination measure, introduce a derived measure of stylistic influence and provide an extension of the measure that can be applied at a group level instead of the individual level.

2.1 Linguistic Style Markers

We measure the linguistic style of a person by their usage of categories of function words that have little semantic meaning, thereby marking style rather than content.

For consistency with prior work, we employed eight of the nine LIWC-derived categories (Pennebaker et al., 2007) deemed to be processed by humans in a generally non-conscious fashion (Ireland et al., 2011). Our eight *markers* are thus: articles, auxiliary verbs, conjunctions, high-frequency adverbs, indefinite pronouns, personal pronouns, prepositions, and quantifiers (451 lexemes total).¹ For experiments where the non-conscious aspect of the markers is not crucial, we also present results for additional LIWC derived stylistic markers: certainty, tentative, discrepancy, inclusive, exclusive, negation, 1st person singular pronouns, 1st person plural pronouns, and 2nd person pro-

¹We discarded negation because it is sparse and seems to carry semantic meaning. (Ireland et al., 2011) also discarded some negations.

Marker	Examples	# of items
Articles	an, the	3
Auxiliary verbs	will, have	144
High frequency adverbs	really, quickly	69
Conjunctions	but, whereas	28
Indefinite pronouns	it, those	46
Personal pronouns	them, her	70
Prepositions	to, with	60
Quantifiers	few, much	89
Certainty	always, never	83
Negation	not, never	57
Discrepancy	should, would	76
Exclusive	without, exclude	17
Inclusive	with, include	18
Tentative	maybe, perhaps	155
1st person singular pronouns	I, me	12
1st person plural pronouns	we, us	12
2nd person pronouns	you, your	20

Table 2.1: Stylistic markers derived from LIWC. Markers in the lower part of the table are only used in some of the experiments, where the non-conscious aspect of the markers was not considered crucial.

nouns. A list of all the markers considered together with examples is provided in Table 2.1.

2.2 Coordination Measure

We start by defining the coordination of one person b towards another person a with respect to a specific linguistic style marker m . We want to quantify how much the use of marker class m in an utterance of a 's *triggers* the occurrence of m in b 's *immediate* (meaning next) reply to that utterance. To put it another way, we want to measure how much a 's use of m in an utterance u_1 increases the probability that b will use m in his reply u_2 , where the increase is relative to b 's normal usage of m in conversations with

a. We stress that we are thus looking at a more subtle phenomenon than whether *b* uses articles (say) more overall when talking to *a*: we want to see whether *b* is so influenced by *a* as to change their function-word usage in their very next reply.

We call *b* the *speaker* and *a* the *target* of a conversational exchange ($a : u_1, b : u_2$), since *a* is the target of *b*'s reply when *b* speaks.² We say an utterance *exhibits* *m* if it contains a word from category *m*. Let $\mathcal{E}_{u_1}^m$ be the event that utterance u_1 (spoken to *b*) exhibits *m*; similarly, let $\mathcal{E}_{u_2 \leftrightarrow u_1}^m$ be the event that reply u_2 to u_1 exhibits *m*.

Given a set $S_{a,b}$ of exchanges ($a : u_1, b : u_2$), we define the coordination of *b* towards *a* as:

$$C^m(b, a) = P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m | \mathcal{E}_{u_1}^m) - P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m) \quad (2.1)$$

where the probabilities are estimated over $S_{a,b}$, and where we require that at least one of *a*'s utterances exhibits *m* in order for the first quantity to be defined.³

2.3 Properties Of The Coordination Measure

2.3.1 Bounds

Eqn. (2.1) has several interesting properties. One non-obvious but important and useful characteristic is that it is a function not only of *b*'s behavior, but also of *a*'s, because it can be shown that (2.1) lies in the interval $\left[-\left(1 - P(\mathcal{E}_{u_1}^m)\right), 1 - P(\mathcal{E}_{u_1}^m)\right]$.

To see why *a*'s behavior needs to be taken into account, consider one extreme case:

²In our terminology, the conversation $\langle x: \text{"Hi." } y: \text{"Tired?" } x: \text{"No."} \rangle$ has two exchanges, one initiated by *x*'s "Hi" (in which *y* is the *speaker* and *x* is the *target*), the other by *y*'s "Tired?" (in which *x* is the *speaker* and *y* is the *target*).

³For clarity, in the figures included in this thesis probabilities will be reported as percentages.

where every utterance of a to b exhibits m . Then $C^m(b, a) = 0$ no matter what b does in response, which makes sense because we have no evidence that any (or no) usage of articles by b is done in response to what a does — we don't have any test cases to see what b does when a doesn't employ a marker.

Another extreme case is also illustrative: where a uses m only a few times when speaking to b , and b uses m when and only when a does. Then, $C^m(b, a)$ approaches 1 as $P(\mathcal{E}_{u_1}^m)$ approaches zero. Again, this makes intuitive sense: it is very unlikely that b matching a exactly on the few times a used m is due merely to chance.

2.3.2 Discussion of alternative measures

In the process of developing the C^m coordination measure, we considered a number of reasonable alternatives . One possibility would be to capture coordination by simply comparing the overall of rate of stylistic marker usage between two persons. However, such a comparison would not reveal the immediate echoing phenomenon we are analyzing and instead would measure a style similarity between the two persons which could be explained by other causes. Another possibility would be to define coordination of b to a as the probability of b using a marker in a reply to an utterance of a that also contains that marker: $P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m \mid \mathcal{E}_{u_1}^m)$. While apparently capturing the immediate nature of the phenomenon, this measure would also suffer from confusion with the background style similarity between a and b : according to this measure two persons with more similar style would appear to coordinate more than two persons with less similar style. This observation motivated the existence of the second term in Equation 2.1, which has the role of controlling for the background similarity of the interlocutors.

Comparison with correlation: the importance of asymmetry Another alternative to $C^m(b, a)$ could also be the well-known correlation coefficient. However, correlation fails to capture an important asymmetry⁴. The case where $\mathcal{E}_{u_1}^m$ but $\neg\mathcal{E}_{u_2 \leftrightarrow u_1}^m$ represents a true failure to accommodate; but the case where $\neg\mathcal{E}_{u_1}^m$ but $\mathcal{E}_{u_2 \leftrightarrow u_1}^m$ should not, at least not to the same degree. For example, u_1 may be very short (e.g., “What?”) and thus not contain an article, but we don’t assume that this completely disallows b from using articles in their reply. In other words, we are interested in whether the presence of m acts as a trigger, not in whether u_2 exhibits m if and only if a does, the latter being what correlation detects.⁵

It bears mentioning that since $\mathcal{E}_{u_1}^m$ and $\mathcal{E}_{u_2 \leftrightarrow u_1}^m$ are binary, a simple calculation shows that the covariance⁶ $cov(\mathcal{E}_{u_1}^m, \mathcal{E}_{u_2 \leftrightarrow u_1}^m) = C^m(b, a) \cdot P(\mathcal{E}_{u_1}^m)$. But, the two terms on the right hand side are not independent: raising $P(\mathcal{E}_{u_1}^m)$ could cause $C^m(b, a)$ to decrease by affecting the first term in its definition, $P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m | \mathcal{E}_{u_1}^m)$ (see equation 2.1).

2.4 Stylistic Influence

Building on the asymmetric nature of C^m , the notion of stylistic influence arises naturally:

$$I_{(b,a)}(m) \triangleq C^m(b, a) - C^m(a, b) \quad (2.2)$$

for a given marker m . If $I_{(b,a)}(m) > 0$ we can say that a accommodates more towards b on m than b does towards a .

⁴Other asymmetric measures based on conditional probability of occurrence have been proposed for adaptation within monologues (Church, 2000) and between conversations (Stenchikova and Stent, 2007). Since our focus is different, we control for different factors.

⁵One could also speculate that it is easier for b to (unconsciously) pick up on the presence of m than on its absence.

⁶The covariance of two random variables is their correlation times the product of their standard deviations.

A related concept is coordination symmetry, which is tied to the coordination measure in the following way. Given that b accommodates to a , i.e $C^m(b, a) > 0$, we have

- Symmetry when $C^m(a, b) > 0$,
- Default asymmetry when $C^m(a, b) = 0$,
- Divergent asymmetry when: $C^m(a, b) < 0$

2.5 Coordination Towards A Group

In the context of group conversations, we can extend the definition of coordination (Equation (2.1)) to the coordination of a particular speaker b towards a *group of targets* A by simply modifying the set of exchanges on which the probabilities in Equation (2.1) are estimated. Specifically, given a set $S_{A,b}$ of exchanges ($a : u_1, b : u_2$) involving initial utterances u_1 of various targets $a \in A$ and replies u_2 of b , the coordination of b to the group A is:

$$C^m(b, A) = P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m | \mathcal{E}_{u_1}^m) - P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m), \quad (2.3)$$

but where this time the probabilities are estimated over $S_{A,b}$.

We then define the coordination of one group of people B towards another group A as the average coordination of speakers in B to targets in A :

$$C^m(B, A) = \langle C^m(b, A) \rangle_{b \in B} \quad (2.4)$$

By taking the macro (unweighted) average, our measure will not be dominated by a few active speakers in a dataset.

2.6 Aggregated Measures

It is important to note that in general, coordination is multimodal: it does not necessarily occur simultaneously for all markers (Ferrara, 1991), and speakers may coordinate on some features but diverge on others (Thakerar et al., 1982). Hence, we also use aggregated measures of coordination of B to A to provide an overall picture of the level of coordination between the groups.

Ideally we want to simply compute $C(b, A)$ as the macro-average of $C^m(b, A)$ across different markers m , and then compute $C(B, A)$ the same way as in (2.4). Recall, however, that $C^m(b, A)$ can only be computed if $S_{A,b}$ contains enough exchanges exhibiting m to reliably estimate both probabilities in (2.3), which is not always the case for all people with respect to all markers. For instance, some persons rarely use quantifiers, leaving C^{quant} undefined in those instances.

We accounted for such “missing values” in three different ways, resulting in three aggregated measures:

Aggregated 1 Compute the “ideal” macro-average $C(b, A)$ only for the persons b for whom $C^m(b, A)$ can be computed for all markers; ignore all the others. This reduces the set of persons considered by the aggregated measure, but provides the most direct measure (in the sense that it does not rely on any particular “smoothing” assumptions as the next two aggregated measures do).

Aggregated 2 For each person b , if $C^m(b, A)$ is undefined, we “smooth” it by using the group average $C^m(B, A)$ instead; this measure considers everybody for which we can compute coordination for at least one marker, but assumes that people in a given group share similar coordination behavior.

Aggregated 3 For each person b , we take the average only over the markers for which $C^m(b, A)$ is defined; this is equivalent to assuming that b would have exhibited the same level of coordination for the missing markers as they did with other markers. This aggregation also considers everybody for which we can compute coordination for at least one marker.

CHAPTER 3

COORDINATION IN SOCIAL MEDIA INTERACTIONS

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 (Java et al., 2007) and that around 37% of all tweets are conversational (Ritter et al., 2010). 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 coordination. 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 coordination 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 coordination is robust enough to occur under these new constraints.

Also, with very few exceptions, coordination was only tested in the initial phase of the development of relations between people (i.e., during the acquaintance process) (Giles, 2008). 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.

3.1 Twitter Conversational Dataset

Drawing from Twitter, Ritter et al. (Ritter et al., 2010) built 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 work. 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 coordination framework proposed in this work and detailed in Chapter 2). 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 only 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 coordination 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

¹<http://apiwiki.twitter.com/>

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

for reconstructing conversations employed in (Ritter et al., 2010)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:

	Mean	Median	Min	Max
Number of conversations	98	60	1	1744
Average number of exchanges	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 exchange*, which is defined as two consecutive tweets in a conversation. The two tweets in an exchange are always sent by different users and are not re-tweets. Conversational dataset A contains 2.6 million exchanges and conversational dataset B contains 420,000 exchanges.

3.2 Empirical Validation

Equipped with the probabilistic framework introduced in Chapter 2, here we proceed with an empirical validation of the coordination phenomenon on the Twitter social media conversation data. As previously discussed, this setting is fundamentally different from all other circumstances in which the theory of communication coordination was validated, therefore challenging its robustness.

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

3.2.1 Validation of stylistic coordination

We now proceed to investigate whether the hypothesis of stylistic coordination proposed in the psycholinguistic literature holds in social media conversations. As described in Section 2.2, for each ordered user pair (a, b) and marker class m we can estimate the subtrahend probability $P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m | \mathcal{E}_{u_1}^m)$ and minuend probability $P(\mathcal{E}_{u_2 \leftrightarrow u_1}^m)$ from Equation (2.1), page 8, over the set $S_{a,b}$ of exchanges $(a : u_1, b : u_2)$ in the data; we will refer to these estimations as $\widehat{PL}_{(b,a)}^m$ and $\widehat{PR}_{(b,a)}^m$, respectively. Then, according to Equation 2.1, we can estimate the amount of coordination C^m exhibited in our dataset as the difference between the mean of the set of subtrahend estimations

$$\{\widehat{PL}_{(b,a)}^m | (a, b) \in \text{Pairs}\}$$

and the mean of the minuend estimations

$$\{\widehat{PR}_{(b,a)}^m | (a, b) \in \text{Pairs}\},$$

where Pairs is the set of all ordered pairs of users in the data⁴. Figure 3.1 compares these means — the former is illustrated in red/right, the latter in blue/left — for each marker class. All the differences are statistically significant with a p-value smaller than 0.0001 according to a two-tailed paired t-test⁵ for all marker classes with the exception of the *2nd person pronoun* marker class for which the difference is not statistically significant.

With the results presented here we are able to verify that coordination does indeed hold in large scale, real world conversational settings 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 coordination exhibits in this conversational setting.

⁴We discard all user pairs for which the denominator of any of these two estimations is less than 10.

⁵In 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) .




Figure 3.1: The effect of coordination C^m for each marker class m observed as the difference between the means of $\{\widehat{P}_{L(b,a)} \mid (a,b) \in \text{Pairs}\}$ (dark, left) and $\{\widehat{P}_{R(b,a)} \mid (a,b) \in \text{Pairs}\}$ (light, right); a dark bar greater than a light bar indicates a positive amount of coordination for the respective marker class. All the differences are statistically significant ($p < 0.0001$), except for the *2nd person pronoun* category. The markers are ordered according to the amount of coordination observed. The y-axis values are reported as percentages for clarity.

3.2.2 Stylistic influence and symmetry

Here we seek to understand the role that the concept of stylistic influence (introduced in Section 2.4) 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 coordination between participants in a conversation by verifying whether we can reject the null hypothesis $E[\text{abs}(I_{(b,a)}(m))] = 0$, where the expectation is taken over all

conversing pairs of users (a, b) . Using definition (2.2), this is reduced to rejecting:

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

and further to rejecting:

$$E [\max (C^m(b, a), C^m(a, b))] = E [\min (C^m(b, a), C^m(a, b))]$$

where the first term is the expected coordination of the most accommodating users (where the coordination is always compared within each pair), and can be estimated by the mean of:

$$\{\max (C^m(b, a), C^m(a, b)) \mid (a, b) \in \text{Pairs}\},$$

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

$$\{\min (C^m(b, a), C^m(a, b)) \mid (a, b) \in \text{Pairs}\}.$$




Figure 3.2: The effect of stylistic influence for each marker class m observed as the difference between the means of $\{\min (C^m(b, a), C^m(a, b)) \mid (a, b) \in \text{Pairs}\}$ (red, left) and $\{\max (C^m(b, a), C^m(a, b)) \mid (a, b) \in \text{Pairs}\}$ (blue, right). All the differences are statistically significant ($p < 0.0001$). The markers are shown in decreasing order of the difference. The y-axis values are reported as percentages for clarity.

Using the same method for estimating $C^m(b, a)$ discussed in Section 3.2.1, we reject this hypothesis for all marker classes m (paired t-test with p-value smaller than 0.0001)⁶. Figure 3.2 illustrates the difference between the expected coordination 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 marker 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 marker classes, we turn our attention to the property of symmetry. Figure 3.3 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 2.4.

The conclusions that can be drawn from analyzing these results is that coordination is a much more complex behavior than previously reported in the literature, where it was assumed that only one type of coordination occurs for a given dimension⁷. But as can be observed in Figure 3.3 all three types of coordination occur to a substantial degree. Furthermore, in all previous work on linguistic style coordination, no distinction was made between the type of coordination occurring for each marker. However, our study indicates a clear difference between markers :

- Symmetric coordination is dominant for *1st pron. pl.*, *Discrepancy* and *Indef. pron.*;

⁶The same holds for all the other dimensions except *Fillers* for which the data was insufficient.

⁷Here we refer to any dimension of coordination, like the ones in Table 1.1, not only to linguistic style dimensions.

- Asymmetric coordination (of both types) is dominant in most of the other dimensions;
- Asymmetric diverging coordination is dominant for *2nd person pronoun*.

A potential explanation for the fact that such a complex coordination behavior was not previously observed may be the difference between the Twitter conversational setting and that traditionally used in the literature (discussed at the beginning of this chapter), 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.

It is natural to ask what factors mediate the imbalance observed in the within-pair amount of coordination. How is linguistic coordination affected by the characteristics of the participants and of their relation? In Chapter 4 we are going to discuss how the power differentials within relations affect the level of coordination imbalance. Studying the relation between other participant characteristics and coordination is an interesting future work direction.

3.2.3 Potential alternative explanations

An additional natural question is, how much are these coordination effects due to an immediate triggering effect, as opposed to simply being a by-product of utterances occurring within the same conversation? For instance, could the results be due just to the topic of the conversation?




Figure 3.3: The percentage of accommodating pairs that exhibit each of the three types of coordination: symmetric, default asymmetric and diverging asymmetric.

Immediate vs. within-conversation effects To answer this question requires measuring “coordination” between utterances that are not adjacent, but are still in the same conversation. To this end, we first restricted attention to those conversations in which there were at least five utterances, so that they would have the structure $a_1 b_2 a_3 b_4 a_5 \dots$. We then measure coordination not between adjacent utterances, like a_1 and b_2 , but where we skip an utterance, such as the pair a_1, b_4 or b_2, a_5 . This helps control for topic effects, since b_4 and a_1 are still close and thus fairly likely to be on the same subject.⁸

Figure 3.4 shows that the level of coordination always falls off after the skipped utterance, sometimes dramatically so, thus demonstrating that the level of immediate adaptation effects we see cannot be solely explained by the topic of conversation or other conversation-level effects. These results accord with the findings of (Levelt and Kelter, 1982), where interposing “interfering” questions lowered the chance of a question’s

⁸It is true that they might be on different topics, but in fact even b_2 might be on a different subject from a_1 .

preposition being echoed by the respondent, and (Reitter et al., 2006), where the effects of structural priming were shown to decay quickly with the distance between the priming trigger and the priming target.




Figure 3.4: Immediate vs. within-conversation effects (for conversations with at least 5 utterances). Suppose that we have a conversation $a_1 b_2 a_3 b_4 a_5 \dots$. The lefthand/dark-green bars show the usual coordination measure, which involves the utterance pair a_1 and b_2 . The righthand/mustard-green bars show coordination based on pairs like a_1 and b_4 — utterances in the same conversation, but not adjacent. We see that there is a much stronger triggering effect for immediately adjacent utterances.

Randomization experiments Towards the same end, we also performed randomization experiments in which we shuffled the order of each participant’s utterances in each conversation, while maintaining alternation between speakers. We again observed drop-offs in this randomized condition in comparison to immediate coordination, the main focus of this thesis.

Temporal effect We also examined whether coordination is affected by the time that has passed between the initial utterance and its reply, using the fact that on Twitter we know the exact time of the utterances. For this purpose we constructed two subsets of the datasets: one that involves only “fast” exchanges, i.e., when the gap between the initial utterance and the reply is less than 5 minutes, and one that involves only “slow” exchanges, i.e., when the gap is larger than 3 hours. As illustrated in Figure 3.5, coordination is significantly higher in “fast” replies than in “slow” replies, in agreement with the results of (Reitter et al., 2006). This result constitutes another piece of empirical evidence supporting that the observed coordination effects are not simply a by-product of utterances occurring within the same conversation (or on the same topic), in which case there would be no reason to observe differences between “fast” and “slow” exchanges.




Figure 3.5: Comparison between coordination levels in “fast” exchanges (exchanges where the reply occurs less than 5 minutes after the initial utterance, dark bars) and “slow” exchanges (exchanges where the reply occurs more than 3 hours after the initial utterance, light bars). We see that there is a stronger triggering effect for fast replies.

Phrase repetition Another alternative explanation is that the coordination effects we are observing are simply due to conscious repetitions inherent to common conversational structure. For example, in the exchange:

Bob: Have you finished **the dissertation**.

Mary: I didn't finish **the dissertation**.

the echoing of the article “the” is not explained by non-conscious stylistic echoing, but simply a result of the (conscious) repetition of the bigram “the dog”. The temporal effect discussed above already discards this phenomenon as the sole explanation of the observed coordination effects, since this alternative explanation would not account for the difference between coordination level in “fast” and “slow” exchanges. Furthermore, we still observe coordination in a dataset where exchanges containing at least one common bigram in the reply and initial utterance were discarded (see also Footnote 8, page 36).

CHAPTER 4

COORDINATION AND POWER DIFFERENCES

In this chapter we show how variations in linguistic style coordination can provide information about power differences within social groups. Our focus here is on domains in which groups engage in goal-oriented discussions — situations where people interact, not necessarily collaboratively, in order to accomplish tasks or settle on choices. An important characteristic of such discussions is that the participants are invested in the issues at hand, so that their dialogs are not simply “idle chat”, but consequential: the outcome matters. Examples include conversations among wiki editors or open-source teams regarding modifications; debates within conference program committees on which papers to accept; and discussions in legal hearings, where opposing sides compete to persuade a judge or jury.

Power differences among the participants constitute a crucial force in all these settings. Sometimes these power differences are embodied in formal roles, such as that of a judge or a program chair. Sometimes they are based on more informal differences in the respect or authority commanded by individuals within the group. And sometimes they are more situational: x may have power over y in a given situation because y needs something that x can choose to provide or not.

We find that differences in the level of language coordination consistently reveal both of these types of power differences in two very different settings. Specifically, we will present the following results.

1. In general, people with low power exhibit greater language coordination than people with high power.

2. Conversely, people coordinate more with interlocutors who have higher power than with those who have lower power.
3. When a person undergoes a change in status, their coordination behavior changes, and so does the coordination behavior of people talking to them.
4. When an individual is trying to convince someone who holds an opposing view, this creates a form of dependence and hence a power deficit in the sense of exchange theory; we find increased levels of language coordination in such cases.
5. The relation between status level and the extent of language coordination transfers across domains, and is a reliable cross-domain feature for status prediction.

These results suggest clear potential applications to the analysis of on-line social groups. In particular, they could provide methods for identifying power differences and levels of status in on-line settings where one has only the text content of social interactions, rather than explicit markers of status or explicitly annotated links. Similarly, they could also provide a means of analyzing conversations between users of a social media platform so as to determine the power balance or levels of relative status in their relationship. In all such uses, the methods do not require domain-specific knowledge of the on-line application being analyzed. We also note that the role of features internal to the content can be crucial in some of these settings, since it has been observed that message frequency and message volume do not necessarily suffice to determine relative status. As Rowe et al. state (Rowe et al., 2007), “As we move down the corporate ladder, the conversational flows of dissimilar employees can in fact be quite similar.” Indeed, it is easy to think of contexts where dominant individuals consume a lot of the conversational bandwidth, and others where, contrariwise, low-status individual take up most of the airtime with their advocacy toward higher-status participants.

There is something striking about the fact that the content features being employed are properties of language that tend to escape conscious attention. The phenomena we find in the text content are consistent and significant, but they are not effects one notices in reading or listening to the interactions; in essence, they operate on levels that only show up when you use computational methods to explicitly tune in to them. Moreover, since our methods are based on function words, it means one can apply them to language samples from which the content words have been redacted, raising intriguing implications for compact representations and user privacy.

4.1 Linguistic Coordination And Power

We can apply communication coordination theory (Giles, 2008; Giles et al., 1991; Natale, 1975; Street and Giles, 1982), an influential line of research in sociolinguistics, to our investigations because the theory implies the following principle:

Principle \mathcal{P} . Linguistic coordination is a function of the power differential between the speaker and the target: the lower the power of the speaker relative to that of the target, the more she coordinates (and vice versa, the higher the relative power of the speaker, the less she coordinates).

Here and throughout, *speaker* refers to the person producing the reply in an exchange, and *target* refers to the person initiating the exchange (and thus the target of the speaker's reply). In the context of group conversations, which is the focus of the present work, this principle leads to the following two concrete hypotheses, based on the power of the target and of the speaker, respectively:

\mathcal{P}_{target} : People in general coordinate more towards *high-powered* people than towards *low-powered* people.

$\mathcal{P}_{speaker}$: *High-powered* people coordinate less than *low-powered* people towards their targets.

(Neither hypothesis implies the other because we employ an asymmetric definition of coordination.)

Using the formal definition of coordination between two groups of people introduced in Section 2.5, we formalize these hypotheses as follows. If people in a group G^{high} have more power than people in a group G^{low} , and U is a set of arbitrary people, the power hypotheses can be rewritten as:

$$\mathcal{P}_{target}: C(U, G^{high}) > C(U, G^{low})$$

$$\mathcal{P}_{speaker}: C(G^{high}, U) < C(G^{low}, U)$$

In addition to power imbalance, we hypothesize that personal traits of the participants also influence how much they coordinate:

\mathcal{B} . People have a baseline coordination level, which is determined by personal characteristics (such as their sociability and level of social engagement).

It is worth noting that it is not actually *a priori* obvious that \mathcal{P}_{target} and $\mathcal{P}_{speaker}$ hold at large. First, there are competing theories which postulate that the relation between power and coordination is the reverse of \mathcal{P} , due to a desire of high-status individuals to be understood (Bell, 1984). Second, empirical studies supporting the hypotheses above are, while intriguing, relatively small in scale. For example, (Gregory Jr. and

Webster, 1996) showed that Larry King, the host of a popular talk-show in the U. S., coordinated more in his vocal pitch to his high-status guests (such as then-President Clinton) than to low-status guests. As for *linguistic style* coordination, (Niederhoffer and Pennebaker, 2002) looked at 15 Watergate transcripts involving only four people altogether (Richard Nixon and three of his aides); small numbers of courtroom trials have also been considered (Aronsson et al., 1987; Erickson et al., 1978).

While power might correlate with certain personal traits in a given community, making the distinction between \mathcal{P} and \mathcal{B} difficult, they differ in one important aspect which we will exploit in our study: power can change abruptly — such as when an individual is assigned a new role — while personal traits, in comparison, are more stable over time. As a result, examining the temporal change in coordination level of people who have undergone changes in power can help us isolate the effect of \mathcal{P} from that of \mathcal{B} . In particular, this will help us address the following question: if we do find evidence supporting hypothesis \mathcal{B} , would it be sufficient to explain the data, or will we see power playing a role on top of baseline individual coordination levels?

4.2 Power Relations In Wikipedia And Supreme Court Data

In this section, we describe the two corpora of consequential discussions we used in our studies. The first consists of discussions between editors on Wikipedia; the second consists of transcripts of oral arguments before the United States Supreme Court. Both settings involve power differentials, both through status and dependence, as we will see below. Our Wikipedia corpus is much larger, potentially more representative of online discussions, and allows us to study the effects of changes in power; but the Supreme Court represents a less collaborative situation than Wikipedia (in the Supreme

Court data, there are always explicit opposing sides) and is an instance of an off-line setting. The differences in the two corpora help us focus on general, domain-independent relationships between relative power and linguistic coordination.

We begin by briefly describing the roles and text content of our two domains, and then discuss how we formalize the different kinds of power imbalances within the domains.

The data is publicly available at <http://www.cs.cornell.edu/~cristian/www2012>.

4.2.1 Discussions among Wikipedia editors

Roles and role changes. Wikipedia editors form a close community with salient markers of status. Administrators, commonly known as *admins*, are Wikipedia editors “trusted with access to restricted technical features” such as protecting or deleting pages or blocking other editors¹. In effect, admins have a higher status than other users (*non-admins*) in the Wikipedia community, and editors seem to be well aware of the status and activity history of other editors. Users are promoted to admins through a transparent election process known as requests for adminship², or *RfAs*, where the community decides who will become admins. Since RfAs are well documented and timestamped, not only do we have the current status of editors, we can also extract the exact time when editors underwent role changes from non-admins to admins.

Textual exchanges. Editors on Wikipedia interact on *talk* pages³ to discuss changes to article or project pages. We gathered 240,436 conversational exchanges carried out

¹<http://en.wikipedia.org/wiki/Wikipedia:Administrators>

²http://en.wikipedia.org/wiki/Wikipedia:Requests_for_adminship

³http://en.wikipedia.org/wiki/Wikipedia:Talk_page_guidelines

on the talk pages, where the participants of these (asynchronous) discussions were associated with rich status and social interaction information: status, timestamp of status change if there is one, and activity level on talk pages, which can serve as a proxy of editors' sociability, or how socially inclined they are. In addition, there is a discussion phase during RfAs, where users "give their opinions, ask questions, and make comments" about an open nomination. Candidates can reply to existing posts during this time. We extracted conversations that occurred in RfA discussions, and obtained a total of 32,000 conversational exchanges. Most of our experiments were carried out on the larger dataset extracted from talk pages, unless otherwise noted.

4.2.2 Supreme Court oral arguments

While Wikipedia discussions provide a large-scale dataset with rich meta-information, overall, high-status people and low-status people are collaborating to accomplish a task. Other social hierarchies involve much less collaboration or even explicitly adversarial relationships. Oral arguments before the Supreme Court provide such a setting.

Roles. A full court consists of nine Justices, although occasionally some recuse themselves. In the oral arguments for a case, lawyers for each party have thirty minutes to present their side to the Justices. The Justices may interrupt these presentations with comments or questions, leading to interactions between the lawyers (plus amici curiae, who for our status-based investigations count as lawyers) and Justices. After the oral arguments and subsequent deliberations, cases are decided by majority vote of the Justices. This provides an interesting additional test ground: instead of asynchronous textual exchanges in a social hierarchy working collaboratively, here we have verbal

		Wikipedia	
		higher power	lower power
Status		admins admins	non-admins admins-to-be (before RfAs)
Dependence		diff. vote	same vote
		Supreme Court	
		higher power	lower power
Status		Justices Chief Justices	lawyers Associate Justices
Dependence		unfavorable Justice	favorable Justice

Table 4.1: Power differentials exhibited in the data

exchanges in a social hierarchy where *Justices* decide the final outcome. In addition, conversations here are over topics in a completely different domain.

Transcripts of verbal exchanges. Transcripts of oral arguments in Supreme Court are publicly available⁴. We used a pre-processed version of this dataset described in (Hawes, 2009). We enhanced this dataset with the final votes from the Spaeth Supreme Court database⁵. In total, we have 50,389 verbal exchanges for 204 cases. 11 justices (two of which have little conversational data: Thomas⁶ and Alito) and 311 lawyers are represented in the dataset. 73% of the lawyers only appear in one case, and the maximum number of cases where one lawyer appears is 15. As such, trends identified on this dataset should not be due to idiosyncratic behavior of a few over-represented lawyers.

4.2.3 Power Relations in the Data

Having now surveyed the nature of the two domains, we discuss the different kinds of power relations that they contain. An overview of the following discussion is summarized in Table 4.1.⁷

In our discussion of roles earlier in this section, we have already indicated some of the basic status differences: the distinction between **admins** and **non-admins** on Wikipedia, and the distinction between **Justices** and **lawyers** in the context of the Supreme Court. We can also identify certain finer-grained distinctions, including the distinction between the **Chief Justice** of the Supreme Court (our data overlaps the terms of two different Chief Justices) and the **Associate Justices**. And on Wikipedia, we can also study the behavior over time of users who were promoted to the position of admin — in effect, comparing their behavior as **admins** to their earlier behavior as **admins-to-be**.

Our data also makes it possible to study several instances of power differences based on *dependence*. To begin with, we note the general principle that status and dependence are almost never completely distinct (Thye, 2000), since a person in a high-status role frequently appears in situations where people are dependent on them.

The data, however, offers us opportunities to study forms of dependence where the level of status has been largely controlled for. Key among these are forms of dependence created by the need to convince someone who disagrees with you. If you are advocating a position in a debate with opposing sides leading to an eventual decision (for example,

⁴http://www.supremecourt.gov/oral_arguments/

⁵<http://scdb.wustl.edu/>

⁶In 2011, Justice Thomas marked five terms without speaking in any oral arguments. (Liptak, 2011)

⁷Throughout this chapter we use color coding to indicate the relative power relations relevant for the respective discussion. These colors are simply intended as a helpful mnemonic and can be ignored without any loss of meaning.

a Supreme Court case, or a policy discussion on Wikipedia prior to a vote), then your audience can be roughly divided into two groups: people who would naturally tend to vote in favor of your position, and people who would naturally tend to vote against your position. Principles of exchange theory indicate that in such situations, you are more dependent on the people who would naturally vote against you, and less dependent on the people who would naturally vote for you, since in order to accomplish your goal, you need to effect a more substantial behavior change in the former group (Emerson, 1962; Kotter, 1977; Wolfe and McGinn, 2005). An important further point here is that in our settings, participants can readily anticipate, either through dialogue or advance knowledge, who is “on their side” and who is “on the other side,” and so it makes sense to suppose that they are aware of these dependence relations during the interaction.

Motivated by this, in the Supreme Court data we will compare levels of coordination of lawyers toward **unfavorable Justices** who (eventually) vote against their side and toward **favorable Justices** who (eventually) vote for their side; there is more dependence and hence more of a power difference in the former case. In the Wikipedia data, we will compare levels of coordination of editors with others who **vote the opposite way** and with others who **vote the same way**; here too, there is more dependence and hence more of a power difference in the former case. We should also note the exchange-theoretic principle that a dependence relation affects both sides: *A*’s dependence on *B* is expected not just to affect *A*’s behavior in their interaction, but *B*’s as well.

4.3 Empirical Investigation

Using the concepts and formalism introduced in the previous sections, we can now investigate the relation between linguistic coordination and power differentials in concrete

conversational settings. Specifically, we test whether the principle \mathcal{P} and the hypotheses \mathcal{P}_{target} and $\mathcal{P}_{speaker}$ introduced in Section 4.1 can be empirically confirmed in the two datasets described in Section 4.2. We begin by discussing power differences arising from status in Wikipedia (where our primary status distinction will be admins vs. non-admins) and in the Supreme Court (where our primary status distinction will be Justices vs. lawyers). After this, we consider power differences arising from dependence.

4.3.1 Power from status: Wikipedia

First, communication behavior on Wikipedia provides evidence for hypothesis \mathcal{P}_{target} : users coordinate more toward the (*higher-powered*) **admins** than toward the **non-admins** (Figure 4.1(a)).⁸

In the other direction, however, when comparing **admins** and **non-admins** as speakers, the data provides evidence that is initially at odds with $\mathcal{P}_{speaker}$: as illustrated in Figure 4.1(b), **admins** coordinate to other people *more* than **non-admins** do (while the hypothesis predicted that they would coordinate *less*).⁹ We now explore some of the subtleties underlying this result, showing how it arises as a superposition of two effects.

Personal characteristics: Hypothesis \mathcal{B} . One possible explanation for the inconsistency of our observations with $\mathcal{P}_{speaker}$ is the effect of personal characteristics suggested in Hypothesis \mathcal{B} from Section 4.1. Specifically, admin status was not conferred arbitrar-

⁸The major explanatory factor for these results does not appear to be wholesale repetition of phrases, even short ones. We note, for example, that with respect to the data used for computing conjunction coordination, only 0.7% of the exchanges contain trigram repeats involving conjunctions and only 3.5% contain bigram repeats involving conjunctions; and the difference in coordination levels remains significant when exchanges with such repeats are discarded.

⁹Note that the observations shown in Figure 4.1(a) do not imply those in Figure 4.1(b), nor vice-versa. For example, the trend in Figure 4.1(a) does not change if we restrict the speakers to be only non-admins (or only admins).




Figure 4.1: Status and linguistic coordination: (a) Users coordinate more towards **admins** (*high-powered*) than towards **non-admins** (*low-powered*), supporting hypothesis \mathcal{P}_{target} (indeed, significantly so in aggregate: see later part of this caption). (b) On the other hand, **admins** (*high-powered*) coordinate more than **non-admins** (*low-powered*) when replying to other people, contradicting hypothesis $\mathcal{P}_{speaker}$.

Note on all figures in this chapter: *’s on the x-axis (e.g., “Article*”) in (a)) indicate statistical significance, independent t-test: * = “ $p < 0.05$ ”, ** = “ $p < 0.01$ ”, *** = “ $p < 0.001$ ”. Next to each legend label, in parentheses, are: the number of users for Aggregated 1 (i.e., the users for which we can compute coordination for all markers) and the total number of users for Aggregated 2 and 3 (i.e., the users for which we can compute coordination for at least one marker). “Error bars” do *not* indicate standard error (we already marked statistical significance with stars) but rather give an idea of how coordination values vary via the standard deviation, estimated by bootstrap resampling (Koehn, 2004). The y-axis values are reported as percentages for clarity.

ily on a set of users; rather, admins are those people who sought out this higher status and succeeded in achieving it. It is thus natural to suppose that, as a group, they may have distinguishing individual traits that are reflected in their level of language coordination.




Figure 4.2: Same-status comparisons

Same-status comparisons (supporting a “winner” personality hypothesis): (a) **admins-to-be** coordinate more than those who remain **non-admins** throughout; (b) during adminship elections (RfAs), **admins-to-be** coordinate more than **failed-to-be**.

Fortunately we can extract rich enough data from Wikipedia that it becomes possible, to a significant extent, to separate the effect of status from these individual traits, establishing that both effects play a role. Our separation of these effects is based on the fact that status can change abruptly, while personal characteristics, though mutable, are more stable over time. On Wikipedia, status changes are well documented, as they can occur only through an election process instigated by requests for adminship (RfAs). When we compare the set of **admins-to-be**— future admins before they were promoted via their RfA — with **non-admins**, Figure 4.2(a) shows that the same differences in language coordination were already present in these two populations — hence, they are not an effect of status alone, since they were visible before the former population experienced a status upgrade.

Can we separate the effects of ambition from success? Yes, because we can look at differences in coordination between users who were promoted (**admins-to-be**), and those who went through the RfA process but were denied admin status (**failed-to-be**). Both **admins-to-be** and **failed-to-be** had the ambition to become admins, but only members of the former group succeeded. We investigate coordination differences between these two groups during a period when their adminship ambitions are arguably most salient: during the discussions in each user’s own RfA process. Figure 4.2(b) shows that even in the conversations they had on their RfA pages, the **admins-to-be** were coordinating more to the others than the **failed-to-be**, providing evidence for a strong form of Hypothesis \mathcal{B} .

Revisiting status: Hypothesis $\mathcal{P}'_{speaker}$. We now return to the issue of status, and describe a method of partially controlling for personal characteristics so as to evaluate the following modification of Hypothesis $\mathcal{P}_{speaker}$:

$\mathcal{P}'_{speaker}$. When controlling for personal characteristics, *high-powered* people coordinate less than *low-powered* people.

To study $\mathcal{P}'_{speaker}$, we create two populations for comparison: the interactions of each **admin** before his or her promotion via RfA (i.e., when they were **admins-to-be**), and the interactions of each **admin** after his or her respective promotion. Figure 4.3(a) shows how the resulting comparison confirms $\mathcal{P}'_{speaker}$: **admins-to-be** decrease their level of coordination once they gain power.¹⁰ Interestingly, the reverse seems to be true for **failed-to-be**: after failing in their RfAs — an event that arguably reinforces their failure to achieve high status in the community — they coordinate more (p-value < 0.05).

¹⁰Note that the trend shown in Figure 4.3(a) is maintained when considering the exact same users in both groups (i.e., excluding the users which did not have enough conversations both before and after adminship). Also note that we allow a time buffer of a month after the RfAs between the two sets of conversations we compare.




Figure 4.3: Effect of status change

Effect of status change. (a) **admins-to-be** coordinate less **after they become admins**; (b) Aggregated 1 coordination of the user (as speaker) and, respectively, towards the user (as target) before and after status change occurs through RfA.

In addition, we can employ status change to reinforce \mathcal{P}_{target} in a setting that controls for personal characteristics: we find that users coordinate more to **admins** after promotion than when they were **admins-to-be** (p-value<0.05).

Finally, in Figure 4.3(b), we investigate how quickly the change in status is reflected in the communication behavior of the users involved. In addition to the monotonic changes in coordination levels over time, and in the hypothesized directions, it is interesting to note that the most dramatic change in coordination is visible in the second month after the change in status occurred. This suggests a period of acclimation to the newly gained status, both for the person that undergoes the change and for those witnessing it.




Figure 4.4: Lawyers coordinate more to Justices than conversely.

4.3.2 Power from status: Supreme Court

In the setting of the Supreme Court, status differences are extremely salient and do not suffer from the correlations that added complexity to the study of $\mathcal{P}_{speaker}$ in its original form. Also, conversations during the oral arguments (almost) always are between a Justice and a lawyer. Thus, our basic finding can be expressed succinctly in Figure 4.4, which shows significantly more coordination from lawyers to Justices than vice versa.¹¹

In the Supreme Court setting we can also study finer-grained status distinctions, to see if these too are manifested in language coordination differences. Indeed, in concordance with \mathcal{P}_{target} , we observe that lawyers coordinate significantly more toward the Chief Justice than toward the Associate Justices (p-value<0.01).

¹¹Throughout, we consider each appearance of a given Justice or lawyer in a different case as a separate entity, which allows for different behaviors in different cases and increases the number of datapoints.




Figure 4.5: Dependence and linguistic coordination: (a) lawyers adjust their coordination level according to whether the Justice is **unfavorable** or **favorable**, supporting \mathcal{P}_{target} ; (b) **favorable Justices** coordinate more than **unfavorable Justices** ($\mathcal{P}_{speaker}$).

4.3.3 Power from dependence

As noted in Section 4.2, we can study power differences based on dependence — even for fixed levels of status difference — using the exchange-theoretic principle that the need to convince someone who disagrees with you creates a form of dependence (Emerson, 1962; Kotter, 1977; Wolfe and McGinn, 2005). Moreover, this power difference is predicted to be felt by both sides — the side with lower power and the side with higher power.

In the case of lawyer-Justice interactions, let us define the Justice to be **favorable** to the lawyer if he or she ends up voting on the lawyer’s side, and **unfavorable** if he or she ends up voting against the lawyer’s side. It is well understood that the Justices often come into the case with a general leaning toward one side or the other based on

their judicial philosophy — this has been acknowledged for example in interviews with members of the Court (Scalia, 2009) — and lawyers through their preparation for the case will come in with knowledge of these leanings. Hence it is reasonable to suppose that the **favorable-unfavorable** distinction will be salient to the interaction during oral arguments.

And indeed, Figures 4.5(a) and 4.5(b) show that the power differences created by this form of dependence are reflected in the amount of coordination, in both directions. First, lawyers coordinate more toward **unfavorable Justices** (on whom they are more dependent) than toward **favorable Justices**, in keeping with \mathcal{P}_{target} . Second, **unfavorable Justices** coordinate less toward lawyers than **favorable Justices** do, in keeping with $\mathcal{P}_{speaker}$. Given the formal framework of *impartiality* that characterizes the Justices’s behavior at the Supreme Court, it is intriguing to see the undercurrent of language coordination differences nevertheless hinting at their eventual decision.

We see a similar effect of dependence on coordination in the context of discussions with opposing sides on Wikipedia. During RfAs, one voter may try to change the opinion of voters on the other side who have already cast their vote. (Changing your vote during the RfA process is allowed, and hence there is an incentive to convince voters to consider this.) Users coordinated more when engaging with **users on the opposite side** than with **those who voted the same way** (Figure 4.6). This finding too, via the arguments about opposing sides and dependence, supports the general power-coordination principle \mathcal{P} .




Figure 4.6: Dependence and linguistic coordination: during RfAs users coordinate more when engaging with **users on the opposite side** than with **those who voted the same way**.

4.4 Cross-Domain Analysis

Part of the motivation for studying the relation between coordination and power is that the principles that govern this relation appear to be domain-independent. Here we perform a set of analyses to show that coordination features do generalize across our two domains more effectively than other text-based features for the problem of inferring power. We find that indeed, compared to the other features we consider, they are the only set of features to display any non-trivial generalization.

Our analysis is based on the following prediction task: for a given pair of different-status people x and y who have engaged in conversations with each other, we predict whether x has the higher status. In this setting, a random guess baseline would achieve 50% accuracy. We stress, however, that this prediction task is primarily a means to assess cross-domain generalization, i.e., not as a free-standing task in itself. Indeed, the

best achievable performance on this status-prediction task appears to be quite domain-dependent. In some domains such as the Supreme Court, idiosyncratic cues in text usage (e.g., lawyers begin their sentences with stylized lead-ins, such as “Your honor”, that clearly mark them as lawyers, not Justices) enable almost perfect performance when these cues are available as features. In other domains, such as Wikipedia, an informal evaluation using two human annotators familiar with the domain produced only 70% accuracy (and an inter-annotator agreement of only 80%). Thus, our interest is not in whether coordination features achieve the best within-domain performance, but in whether they are particularly effective at generalizing (as we indeed find them to be).

Experimental setup. Let R_x be x ’s replies to y , and R_y be y ’s replies to x , and $Len(S)$ be the average length of all utterances in the set S . Let \mathcal{F}_{style} be the set of 8 stylistic markers introduced in Section 2.1. We define the following sets of features used as input to an SVM classifier:

- coordination features: binary features indicating, for each $m \in \mathcal{F}_{style}$ as well as for Aggregated 1¹² (as introduced in Section 2.6, page 12), whether x coordinates more to y than y to x on m
- stylistic features: frequency of each marker $m \in \mathcal{F}_{style}$ in R_x and, respectively, in R_y ; also, $Len(R_x)$, $Len(R_y)$. We use this feature set to examine whether style alone is predictive on its own, or whether specifically stylistic *coordination* is key
- bag of words: frequency of each word in R_x , frequency of each word in R_y , L_2 -normalized

For experiments on the Wikipedia data, which we denote as *wiki*, we considered

¹²We only considered pairs of participants for which enough data was available to compute coordination on all stylistic features.

Training corpus	in-domain		cross-domain	
	<i>wiki</i>	<i>court</i>	<i>court</i>	<i>wiki</i>
Test corpus	<i>wiki</i>	<i>court</i>	<i>wiki</i>	<i>court</i>
coordination features (9 altogether)	57.7	70.4	57.1	55.0
stylistic features (18 altogether)	59.2	51.4	50.0	51.9
bag of words (20,000 altogether)	51.4	99.5	45.2	40.1

Table 4.2: Prediction accuracy for SVM’s using various feature sets. Cross-domain results are in the right-hand two columns. Bold = results significantly better than chance.

(*admin, non-admin*) pairs (for conversations occurring after the admins were elected). For the Supreme Court dataset (*court*), we considered (*Justice, lawyer*) pairs¹³.

For *in-domain* experiments, we report average accuracy over cross-validation within the same domain (i.e., training and test corpora are both *wiki* or *court*); for *cross-domain* experiments, we train on one domain and test on the other.

Results. Table 4.2 summarizes the results. We find that coordination features are the only ones to perform statistically significantly better than random guessing in the *cross-domain* settings — the other classifiers simply learn cues that are idiosyncratic to their training data, and fail to generalize. (Note for example that the bag-of-words method picks up on the near-perfect lexical cues marking lawyers in the Supreme Court data, but this method performs worse than random guessing when applied to the other domain.)

Even looking at the *in-domain* tasks — which were not our primary focus here — we find that coordination features are the only ones that perform statistically significantly better than random guessing on *both* datasets.

¹³In order to focus on the conversational exchanges and avoid exchanges in which the lawyers formally introduce their case, we considered only cases where the length difference between the two utterances were fewer than 20 words.

Hypothesis	Wiki	Wiki > 20	S.C.	S.C. > 20
low use of I-words	holds	holds	holds	reverse
high use of we-words	no diff.	no diff.	reverse	no diff.
high use of you-words	no diff.	no diff.	holds	holds

Table 4.3: Testing the pronoun-status hypotheses (Pennebaker, 2011) by comparing the respective relative frequencies in Wikipedia replies, Wikipedia replies longer than 20 words, Supreme Court replies and Supreme Court replies longer than 20 words. “holds” indicates that the hypothesis was confirmed in the respective dataset, “reverse” indicates that the opposite of the hypothesis holds, and “no diff.” indicates that no statistical significant difference in pronoun usage was found.

Personal pronouns as markers of status Psycholinguistic studies (Pennebaker, 2011) suggest that the usage of personal pronouns can also reveal status differences. In particular, it was hypothesized that high status people have a low use of first person singular pronouns (“I-words”), a high use of first person plural pronouns (“we-words”), and a high use of second person pronouns (“you-words”). This motivated us to conduct some preliminary investigations on whether personal pronouns can be used as reliable cross-domain markers of status. Our results, summarized in Table 4.3, indicate that these hypotheses are not entirely supported by the Wikipedia and Supreme Court data.

4.5 Interactions Among Hypotheses \mathcal{P} And \mathcal{B}

In Section 4.3 we saw that the interaction between personal characteristics (which form the basis for Hypothesis \mathcal{B}) and power differentials (which form the basis for Hypothesis \mathcal{P}) can lead to complex effects. Here we consider two cases where this interaction raises interesting issues, and point to open questions in the analysis of coordination.

An individual’s level of social engagement is one type of personal characteristic that interacts with coordination and power. As a simple proxy for social engagement, for

purposes of discussion here, we consider the volume of communication the individual engages in. As we noted at the beginning of this chapter, simple volume measures such as this do not seem to readily yield domain-independent information about power, since they vary considerably across domains — in some domains the powerful people talk a lot, and in others they talk relatively little. For example, when people are promoted to **admin** status, their volume of communication goes up while (as we have seen) their level of coordination goes down. On the other hand, **lawyers** talk more than **Justices** in the Supreme Court data, and (again as we have seen) they also coordinate more in the **lawyer-Justice** interactions.

However, if we restrict attention to a fixed sub-population within a given domain, there are interesting connections between coordination and volume that suggest further questions. In particular, on Wikipedia we consider the number of replies posted by a user on talk-pages as a measure of communication volume, and hence a proxy for their level of social engagement on the site. We compared users in the top 1/3 of the sorted order by communication volume with users in the bottom 1/3, finding that users with higher numbers of replies are more likely to coordinate to others (p-value<0.05). We observed the same effect when we compared the communication volumes of users with the same status: among admins, users with more communication are also more likely to coordinate, and the same trend holds among non-admins. Similar effects also hold for other measures of communication volume. Again, we note that other domains (such as the Supreme Court) show an inverse relation between volume and coordination in the communication transcripts, and so it is an interesting question to identify the principles that determine how this relationship plays out in different settings.

We also consider a second basic example that raises an interesting challenge for distinguishing between Hypotheses \mathcal{P} and \mathcal{B} : the effect of gender on coordination, us-

ing the fact that gender information is available for participants in the Supreme Court dataset.




Figure 4.7: Gender differences

The main finding here, in Figure 4.7, is that overall female lawyers coordinate more than male lawyers when talking to Justices, , and correspondingly, Justices coordinate more towards male lawyers than towards female lawyers. Given the extensive history of work exploring culturally-embedded status and power differences based on gender (Berger et al., 1972; Strodtbeck and Mann, 1956), one interpretation of this finding is directly in terms of Hypothesis \mathcal{P} . However, since it is also potentially related to theories of gender-based communication differences (Herring, 2003) and even gender-based language adaptation differences (Otterbacher and Hemphill, 2012), the question of separating Hypotheses \mathcal{P} and \mathcal{B} becomes challenging here. We think it is a promising possibility that language coordination effects may be able to serve as a lens through which

to measure many similar kinds of distinctions in both on-line and off-line conversational settings.

CHAPTER 5

CAUSAL MECHANISM

People don't consciously track function words (Levelt and Kelter, 1982; Petten and Kutas, 1991; Segalowitz and Lane, 2004) — it's not easy to answer the question, "how many prepositions were there in the sentence I just said?". Therefore, it is quite striking that humans nonetheless instantly adapt to each other's function-word rates. Indeed, there is active debate regarding what mechanisms cause nonconscious coordination (Branigan et al., 2010; Ireland et al., 2011).

One line of thought is that coordination represents a social strategy¹ whose aim is to gain the other's social approval (Giles, 2008; Street and Giles, 1982) or enhance the other's comprehension (Bortfeld and Brennan, 1997; Clark, 1996).² This hypothesis is supported by studies showing that coordination is affected by a number of social factors, including relative social status (Gregory Jr. and Webster, 1996; Natale, 1975; Thakerar et al., 1982) and gender role (Bilous and Krauss, 1988; Ireland and Pennebaker, 2010; Namy et al., 2002). The relation between coordination and power differentials revealed in Chapter 4 also supports this hypothesis.

But an important question is whether the adaptation mechanism has become so deeply embedded in the language-generation process as to have transformed into a reflex not requiring any social triggering.³ Indeed, it has been argued that unconscious mimicry is partly innate (Chartrand and Bargh, 1999), perhaps due to evolutionary pressure to foster relationships (Lakin et al., 2003).

¹In fact, social signaling may also be the evolutionary cause of chameleons' color-changing ability (Stuart-Fox et al., 2008).

²For the purpose of our discussion, we are conflating social-approval and audience-design hypotheses under the category of *social strategy*.

³This hypothesis relates to characterizations of alignment as an unmediated mechanism (Pickering and Garrod, 2004).

To answer this question, we take a radical approach: we consider a setting in which the persons *generating* the coordinating dialog are different from those *engaged* in the dialog (and standing to reap the social benefits) — imagined conversations, specifically, scripted movie dialogs.

Life is beautiful, but cinema is paradise A priori, it is not clear that movie conversations would exhibit coordination. Dialogs between movie characters are not truthful representations of real-life conversations. They often are “too carefully polished, too rhythmically balanced, too self-consciously artful” (Kozloff, 2000), due to practical and artistic constraints and scriptwriting practice (McKee, 1999). For example, mundane phenomena such as stuttering and word repetitions are generally nonexistent on the big screen. Moreover, writers have many goals to accomplish, including the need to advance the plot, reveal character, make jokes as funny as possible, and so on, all incurring a cognitive load.

So, the question arises: do scripted movie dialogs, in spite of this quasi-artificiality and the aforementioned generation/engagement gap, exhibit the real-life phenomenon of stylistic coordination? When imagining dialogs, do scriptwriters (nonconsciously⁴) adjust the respondent’s replies to echo the initiator’s use of articles, prepositions, and other apparently minor aspects of lexical choice? According to our results, this is indeed the case, which has fascinating implications.

First, this provides evidence that coordination, assumed to be driven by social motivations, has become so deeply embedded into our ideas of what conversations “sound like” that the phenomenon occurs even when the person generating the dialog is not the recipient of the social benefits.⁵

⁴The phenomenon of real-life language coordination is not widely known among screenplay authors (Beth F. Milles, professor of acting and directing, personal communication).

⁵Although some writers may perhaps imagine themselves “in the shoes” of the recipients, recall that authors generally don’t include in their scripts the repetitions and ungrammaticalities of “real-life” speech.

Second, movies can be seen as a controlled environment in which preconceptions about the relation between communication patterns and the social features of the participants can be studied. This gives us the opportunity to understand how people (scriptwriters) nonconsciously *expect* coordination to relate to factors such as gender, status and relation type. Are female characters thought to accommodate more to male characters than vice-versa?

Furthermore, movie scripts constitute a corpus that is especially convenient because meta-features like gender can be more or less readily obtained.

Contributions We check for coordination in a corpus of roughly 250,000 conversational exchanges from movie scripts and find a statistically significant coordination effect for all stylistic markers considered. We thereby provide evidence that language coordination is so implanted within our conception of conversational behavior that, even if such coordination is socially motivated, it is exhibited even when the person generating the language in question is not receiving any of the presumed social advantages.

We also study the effects of gender, narrative importance, and hostility. Intriguingly, we find that these factors indeed “affect” movie characters’ linguistic behavior; since the characters aren’t real, and control of stylistic lexical choice is largely nonconscious, the effects of these factors can only be springing from patterns existing in the scriptwriters’ minds.

5.1 Movie Dialogs Corpus

To address the questions raised in the introduction, we created a large set of imagined conversations, starting from movie scripts crawled from various sites.⁶ Metadata for conversation analysis and duplicate-script detection involved mostly-automatic matching of movie scripts with the IMDB movie database; clean-up resulted in 617 unique titles tagged with genre, release year, cast lists, and IMDB information. We then extracted 220,579 conversational exchanges between pairs of characters engaging in at least 5 exchanges, and automatically matched these characters to IMDB to retrieve gender (as indicated by the designations “actor” or “actress”) and/or billing-position information when possible (≈ 9000 characters, ≈ 3000 gender-identified and ≈ 3000 billing-positioned). The latter feature serves as a proxy for narrative importance: the higher up in the credits, the more important the character tends to be in the film.

To the best of our knowledge, this is the largest dataset of (metadata-rich) imaginary conversations to date.

5.2 Experimental Results

5.2.1 Coordination exists in fictional dialogs

For each ordered pair of characters (a, b) and for each feature family m , we estimate the amount of coordination given by equation (2.1) in the same manner as in Section 3.2, page 16. Fig. 5.1 compares the average values of these subtrahend and minuend means

⁶The source of these scripts and more detail about the corpus are given in the README associated with the Cornell movie-dialogs corpus, available at <http://www.cs.cornell.edu/~cristian/movies>.




Figure 5.1: Implicit depiction of coordination for each marker m , illustrated as the difference between the means of $\{\widehat{PL}_{(b,a)} \mid (a, b) \in \text{Pairs}\}$ (right/light-blue bars) and $\{\widehat{PR}_{(b,a)} \mid (a, b) \in \text{Pairs}\}$ (left/dark-blue bars). (This implicit representation allows one to see the magnitude of the two components making up our definition of coordination.) The markers are ordered by decreasing coordination. All differences are statistically significant (paired t-test). In all figures in chapter, error bars represent standard error, estimated via bootstrap resampling (Koehn, 2004). (Here, the error bars, in red, are very tight.)

(as a way of putting coordination values into context), showing positive differences for all of the considered families of features (statistically significant, paired t-test $p < 0.001$); this demonstrates that movie characters do indeed converge to each other’s linguistic style on all considered trigger families.⁷

Movies vs. Twitter One can ask how our results on movie dialogs correspond to those for real-life conversations. To study this, we utilize the results presented in Chapter 3

⁷We obtained the same qualitative results when measuring coordination via the correlation coefficient, doing so for the sake of comparability with prior work (Niederhoffer and Pennebaker, 2002; Taylor and Thomas, 2008).




Figure 5.2: Coordination in Twitter conversations (left bars) vs. coordination in movie dialogs (right bars; corresponds to the difference between the two respective bars in Fig. 5.1) for each trigger family. The trigger families are ordered by decreasing coordination in Twitter.

on Twitter as data on real conversational exchanges. Figure 5.2 depicts the comparison, revealing two interesting effects. First, Twitter users coordinate more than movie characters on all the trigger families we considered, which does show that the coordination effect is stronger in actual interchanges. On the other hand, from the perspective of potentially using imagined dialogs as proxies for real ones, it is intriguing to see that there is generally a correspondence between how much coordination occurs in real dialogs for a given feature family and how much coordination occurs for that feature in imagined dialogs, although conjunctions and articles show a bit less coordination in fictional exchanges than this pattern would suggest.




Figure 5.3: Relation between *Article* coordination and imagined gender. (a) compares cases when the **initiator** and **respondent** are **Male** or **Female**; (b) compares types of gendered **initiator-respondent** relations: **Male-Male**, **Female-Male**, **Male-Female**, **Female-Female**. For comparison, the **All** bars represents the general *Article* coordination (illustrated in Fig. 5.1 as the difference between the two respective bars).

5.2.2 Self-coordination

Could our results be explained entirely by the author converging to their own self, given that self-alignment has been documented (Pickering and Garrod, 2004; Reitter et al., 2006)? If that were the case, then the *characters* that the author is writing about should converge to themselves no more than they converge to different characters. But we ran experiments showing that this is not the case, thus invalidating this alternative hypothesis: in fact, characters converge to themselves much more than they converge to other characters (Figure 5.4(a))

There is a notable qualitative distinction between imagined conversations and real conversations (on Twitter) with respect to self-coordination: in real conversations self-coordination has a considerably smaller amplitude than coordination, while the opposite

is true in movies (compare Figure 5.4(a) with Figure 5.4(b)). This distinction could potentially be employed in detecting forged conversations, in a way similar to that in which Benford’s law is employed for detecting forged financial reports (Cho and Gaines, 2007). 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.




Figure 5.4: Coordination (dark bars) vs self-coordination (light bars) in movies conversations (a) and in Twitter conversations (b).

5.2.3 Coordination and imagined relation

We now analyze how coordination patterns vary with the type of relationship between the (imagined) participants. Note that, given the multimodal character of coordination, treating each marker class separately is the most appropriate way to proceed, since in

past work, for the same experimental factor (e.g., gender), different features converge differently (Bilous and Krauss, 1988). For clarity of exposition, we discuss in detail only the results for the *Articles* marker class; but the results for all marker classes are summarized in Fig. 5.7, discussed later.

Imagined gender Fig. 5.3(a) shows how coordination on article usage depends on the gender of the initiator and respondent. Females are more influential than males: movie characters of either gender accommodate more to female characters than to male characters (compare the **Female initiator** bar with the **Male initiator** bar, statistically significant, independent t-test, $p < 0.05$). Also, female characters seem to accommodate slightly more to other characters than male characters do (though not statistically significantly so in our data).

We also compare the amount of coordination between all the possible types of gendered initiator-respondent pairs involved (Fig. 5.3(b)). One can observe, for example, that male characters adapt less in same-gender situations (**Male-Male** conversations) than in mixed-gender situations (**Female initiator-Male** respondent), while the opposite is true for female characters (**Female-Female** vs. **Male-Female**).

Interpreting these results lies beyond the scope of this work. We note that these results could be a correlate of many factors, such as the roles that male and female characters are typically assigned in movie scripts.

Narrative importance Does the relative importance bestowed by the scriptwriter to the characters affect the amount of linguistic coordination he or she (nonconsciously) embeds in their dialogs? Fig. 5.5 shows that, on average, the lead character converges

to the second-billed character more than vice-versa (compare left bar in **1st resp.** group with left bar in **2nd resp.** group).

One possible confounding factor is that there is significant gender imbalance in the data (82% of all lead characters are males, versus only 51% of the secondary characters). Could the observed difference be a direct consequence of the relation between gender and coordination discussed above? The answer is no: the same qualitative observation holds if we restrict our analysis to same-gender pairs (compare the righthand bars in each group in Fig. 5.5⁸).

It would be interesting to see whether these results could be brought to bear on previous results regarding the relationship between social status and coordination, but such interpretation lies beyond the scope of this work, since the connection between billing order and social status is not straightforward.

Quarreling The level of contention in conversations has also been shown to be related to the amount of coordination (Giles, 2008; Niederhoffer and Pennebaker, 2002; Taylor and Thomas, 2008). To test whether this tendency holds in the case of imagined conversations, as a small pilot study, we manually classified the conversations between 24 main pairs of characters from romantic comedies⁹ as: *quarreling*, *some quarreling* and *no quarreling*. Although the experiment was too small in scale to provide statistical significance, the results (Fig. 5.6) suggest that indeed the level of coordination is affected by the presence of controversy: *quarreling* exhibited considerably more coordination

⁸Figure 5.5 also shows that our coordination measure does achieve negative values in practice, indicating divergence. Divergence is a rather common phenomenon which deserves attention in future work; see Chapter 3 for an account.

⁹We chose the romantic comedy genre since it is often characterized by some level of contention between the two people in the main couple.




Figure 5.5: Comparison of the coordination of first-billed (lead) characters to second-billed characters (left bar in **1st resp.** group) to that of second-billed characters to leads (left bar in **2nd resp.** group); righthand bars (dark green) in each group show results for Male-Male pairs only.

for articles than the other categories (the same holds for personal and indefinite pronouns; see Fig. 5.7(d)). Interestingly, the reverse is true for adverbs; there, we observe divergence for contentious conversations and coordination for non-contentious conversations (Fig. 5.7(d)). This corresponds to Niederhoffer and Pennebaker’s (Niederhoffer and Pennebaker, 2002) observations made on real conversations in their study of the Watergate transcripts: when the relationship between the two deteriorated, Richard Nixon converged more to John Dean on articles, but diverged on other features.¹⁰

Results for the other features Our results above suggest some intriguing interplay between coordination and gender, status, and level of hostility in imagined dialogs, which may shed light on how people (scriptwriters) nonconsciously *expect* coordination to relate to such factors. (Interpreting these sometimes apparently counterintuitive findings is beyond the scope of this work, but represents a fascinating direction for future work.)

¹⁰Adverbs were not included in their study.






Figure 5.6: Relation between contention and coordination. The third bar combines *quarreling* and *some quarreling* to ameliorate data sparsity. For comparison, **Rom. com.** shows coordination calculated on all the conversations of the 24 romantic-comedy pairs considered in this experiment.


Fig. 5.7 shows how the nature of these relations depends on the marker class considered. The variation among families is in line with the previous empirical results on the multimodality of coordination in real conversations, as discussed in Section 2.6.




(a) **F resp.** minus **M resp.**



(b) **F init.** minus **M init.**



(c) **1st resp.** minus **2nd resp.**



(d) **Quarrel** minus **No quarrel**

Figure 5.7: Summary of the relation between coordination and imagined gender (a and b), billing order (c), and quarreling (d). The bars represent the *difference* between the coordination observed in the respective cases; e.g., the **Article** (red) bar in (a) represents the difference between the **F resp.** and the **M resp.** bars in Fig. 5.3(a). In each plot, the trigger families are sorted according to the respective difference, but the color assigned to each family is consistent across plots. The scale of (d) differs from the others.

CHAPTER 6

FURTHER RELATED WORK

6.1 Linguistic Style And Human Characteristics

Using stylistic (i.e., non-topical) elements like articles and prepositions to characterize the utterer in some way has a long history. Linguistic style was shown to be crucial in the area of authorship attribution and in forensic linguistics (for an overview see (Juola, 2008)). 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. (Holmes, 1994; Yule, 1939)). Since Mosteller and Wallace’s seminal work on the Federalist Papers (Mosteller and Wallace, 1963), 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 (Herring and Paolillo, 2006; Koppel et al., 2002; Mukherjee and Liu, 2010) and personality type detection (Argamon et al., 2005; Mairesse et al., 2007; Oberlander and Gill, 2006), identification of interactional style (Jurafsky et al., 2009; Ranganath et al., 2009), and recognizing deceptive language (Hancock et al., 2008; Mihalcea and Strapparava, 2009; Ott et al., 2011).

6.2 Language In Social Media

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 anal-

ysis of the content of posts as well as structural relations among posters. In one early study combining these two dimensions of analysis, Paolillo (Paolillo, 2001) 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 (?) to various language features including intimacy words, positive and negative emotions. Eisenstein et al. (Eisenstein et al., 2010) investigated the role geographic variation of language has in Twitter and Kiciman (Kiciman, 2010) examined the extent to which differences in language models of Twitter posts (as measured by perplexity) were related to metadata associated with the senders. Ramage et al. (Ramage et al., 2010) 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 personal 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 (Duan et al., 2010) which employs twitter specific features in conjunction to textual content.

Other work focused on the structural properties of conversations in social media. Honeycutt and Herring (Honeycutt and Herring, 2009) 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. (Ritter et al., 2010) 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 in-

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.

6.3 Power, Structural Features And Language

There has been extensive work on using structural features, rather than language, to infer notions of “importance” in networks, both in the literature on social networks (Wasserman and Faust, 1994) and on the Web (Chakrabarti, 2003). Recent work has also studied the inference of status from on-line social network features (Guha et al., 2004; Leskovec et al., 2010b).

The relation between linguistic coordination¹ and status has been studied in an interesting but small-scale study (Niederhoffer and Pennebaker, 2002) which looked at 15 Watergate transcripts involving only four people altogether (as such, it did not explore the bulk of the phenomena we identify here). Additionally, researchers have used text features other than linguistic coordination to identify status differences (Bramsen et al., 2011; Diehl et al., 2007; McCallum et al., 2007; Scholand et al., 2010); in contrast with our work these methods picked up situation-specific cues, such as the word “termination” for the Enron corporate-email corpus (Diehl et al., 2007), which are unlikely to generalize across contexts.

¹For brevity, we exclude examinations of the effects of status on other types of coordination, such as pitch and vocal features, which are absent from textually-manifested discussions (see (Giles, 2008) for a survey) or on related phenomena such as content matching (Aronsson et al., 1987).

6.4 Collaborative Communities

An extensive amount of work has contributed to a better understanding of the rules that govern collaborative behavior in online communities. Wikipedia was used as a testbed for studying user interaction at large (Billings and Watts, 2010; Laniado et al., 2011; Lu et al., 2010; Taraborelli and Ciampaglia, 2010; Viégas et al., 2007) and the promotion process in such communities (Burke and Kraut, 2008; Leskovec et al., 2010a). Reviewer behavior and incentives to participate in the collaborative process were studied in the context of commercial review sites (Bryant et al., 2005; Gilbert and Karahalios, 2009; Lu et al., 2010; Wu and Huberman, 2010).

6.5 Imagined Conversations

There has been work in the NLP community applying computational techniques to fiction, scripts, and other types of text containing imagined conversations. For example, one recent project identifies conversational networks in novels, with the goal of evaluating various literary theories (Elson and McKeown, 2010; Elson et al., 2010). Movie scripts were used as word-sense-disambiguation evaluation data as part of an effort to generate computer animation from the scripts (Ye and Baldwin, 2006). Movie scripts were also employed to study how linguistic features can affect the memorability of phrases (Danescu-Niculescu-Mizil et al., 2012a). Computational studies of poetry included an analysis of the relationship between pronunciation and network structure (Sonderegger, 2010), classification of professional and amateur poetry (Kao and Jurafsky, 2012) and quantification of stylistic segmentation (Brooke et al., 2012). Part-of-speech frequencies were computed for imaginative writing in the British National Corpus, finding a typology gradient progressing from conversation to imaginative writing

(e.g., novels) to task-oriented speech to informative writing (Rayson et al., 2001). The data analyzed by (Oberlander and Gill, 2006) consisted of emails that participants were instructed to write by imagining that they were going to update a good friend on their current goings-on.

CHAPTER 7

SUMMARY AND FUTURE DIRECTIONS

This thesis proposes a computational approach to linguistic style coordination, enabling a new understanding of this intriguing aspect of conversational behavior. A new computational framework is developed and applied to four conversational settings — Twitter conversations, discussions between Wikipedia editors, oral arguments in the U.S. Supreme Court and movie script dialogs — with the goal of exploring different dimensions of the phenomenon.

We show for the first time that coordination emerges in social media and by analyzing the phenomenon in this large scale setting we reveal important properties which were never before observed. Building on the resulting insights, we show that linguistic coordination provides information about the power differences between individuals taking part in the social interactions within a group. While traditional methods for inferring social structure at large scales have used network properties, our findings thus suggest methods for taking information from the textual content of social interactions and bringing it to bear on these types of problems.

We also provide some insight into the causal mechanism behind coordination, a topic that has generated substantial scrutiny and debate for over 40 years (Branigan et al., 2010; Ireland et al., 2011). These results, along with (Elson and McKeown, 2010), advocate for the value of fictional sources in the study of linguistic and social phenomena.

This line of research has a number of natural extensions. One question we have not addressed is the issue of long-term adaptation: can we measure the changes of linguistic style over a longer period of time, as the relation between participants develops? One of

the main challenges here is to design a method for measuring and clustering patterns of linguistic change over time.

Another set of questions is to further understand the types of social distinctions that are manifested by differences in coordination; while we have seen that multiple forms of power can be observed this way, there could well be other properties that can be exposed as well. More broadly, our work makes clear that language use contains subtle features that reveal latent social information, and identifying new classes of such features promises to extend our understanding of group interaction.

BIBLIOGRAPHY

- The Getaway. Written by Walter Hill, directed by Sam Peckinpah, 1972.
- Shlomo Argamon, Sushant Dhawle, Moshe Koppel, and James W. Pennebaker. Lexical predictors of personality type. *Proceedings of the 2005 Joint Annual Meeting of the Interface and the Classification Society of North America*, 2005.
- Karin Aronsson, Linda Jönsson, and Per Linell. The courtroom hearing as a middle ground: Speech accommodation by lawyers and defendants. *J. Lang. and Social Psych.*, 6(2):99–115, 1987.
- Allan Bell. Language style as audience design. *Language in Society*, 13(2):145–204, 1984.
- Joseph Berger, Bernard P. Cohen, and Morris Zelditch Jr. Status characteristics and social interaction. *American Sociological Review*, 37(3):241–255, 1972.
- Matt Billings and Leon A Watts. Understanding dispute resolution online: Using text to reflect personal and substantive issues in conflict. In *CHI*, pages 1447–1456, 2010.
- Frances Bilous and Robert Krauss. Dominance and accommodation in the conversational behavior of same- and mixed-gender dyads. *Language and Communication*, 8: 183–194, 1988.
- J. Kathryn Bock. Syntactic persistence in language production. *Cognitive Psychology*, 18(3):355 – 387, 1986.
- Heather Bortfeld and Susan E. Brennan. Use and acquisition of idiomatic expressions in referring by native and non-native speakers. *Discourse Processes*, 23(2):119–147, 1997.

- Philip Bramsen, Martha Escobar-Molana, Ami Patel, and Rafael Alonso. Extracting social power relationships from natural language. In *ACL HLT*, 2011.
- Holly P. Branigan, Martin J. Pickering, Jamie Pearson, and Janet F. McLean. Linguistic alignment between people and computers. *Journal of Pragmatics*, 42(9):2355–2368, 2010.
- Julian Brooke, Adam Hammond, and Graeme Hirst. Unsupervised Stylistic Segmentation of Poetry with Change Curves and Extrinsic Features. In *Proceedings of the NAACL Workshop on Computational Linguistics for Literature*, page 26, 2012.
- Susan L Bryant, Andrea Forte, and Amy Bruckman. Becoming Wikipedian: transformation of participation in a collaborative online encyclopedia. In *GROUP*, pages 1–10, 2005.
- Moira Burke and Robert Kraut. Mopping up: Modeling Wikipedia promotion decisions. In *CSCW*, pages 27–36, 2008.
- Soumen Chakrabarti. *Mining the Web: discovering knowledge from hypertext data*. Morgan Kaufmann, 2003.
- Tanya L. Chartrand and John A. Bargh. The chameleon effect: The perception-behavior link and social interaction. *J. Personality and Social Psychology*, 76(6):893–910, 1999.
- Tanya L. Chartrand and Rick van Baaren. Chapter 5: Human mimicry. In Mark P. Zanna, editor, *Advances in Experimental Social Psychology*, volume 41, pages 219–274. Academic Press, 2009.
- Wendy K. Tam Cho and Brian J. Gaines. Breaking the (benford) law: Statistical fraud detection in campaign finance. *The American Statistician*, 61(3):218–223, 2007.

- Kenneth W. Church. Empirical estimates of adaptation: the chance of two noriegas is closer to $p/2$ than p^2 . In *Proceedings of COLING*, pages 180–186, 2000.
- Herbert H. Clark. *Using language*. Cambridge University Press, second edition, 1996.
- William S. Condon and William D. Ogston. A segmentation of behavior. *Journal of psychiatric research*, 1973.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *ACL Cognitive Modeling and Computational Linguistics Workshop*, 2011.
- Cristian Danescu-Niculescu-Mizil, Michael Gamon, and Susan Dumais. Mark my words! Linguistic style accommodation in social media. In *Proceedings of WWW*, pages 745–754, 2011.
- Cristian Danescu-Niculescu-Mizil, Justin Cheng, Jon Kleinberg, and Lillian Lee. You had me at hello: How phrasing affects memorability. In *Proceedings of ACL*, pages 892–901, 2012a.
- Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang, and Jon Kleinberg. Echoes of power: Language effects and power differences in social interaction. In *Proceedings of WWW*, pages 699–708, 2012b.
- Valerian J. Derlega, Marian S. Harris, and Alan L. Chaikin. Self-disclosure reciprocity, liking and the deviant. *Journal of Experimental Social Psychology*, 9(4):277–284, 1973.
- Christopher P. Diehl, Galileo Namata, and Lise Getoor. Relationship identification for social network discovery. In *Proceedings of the AAAI Workshop on Enhanced Messaging*, volume 22, pages 546–552, 2007.

- John W Du Bois. Towards a dialogic syntax. Unpublished manuscript, 2001.
- Yajuan Duan, Long Jiang, Tao Qin, Ming Zhou, and Hheung-Yeung Shum. An empirical study on learning to rank of tweets. In *Proceedings of the 23rd International Conference on Computational Linguistics*, pages 295–303, 2010.
- Jacob Eisenstein, B O’Connor, Noah A Smith, and Eric P Xing. A latent variable model for geographic lexical variation. *EMNLP*, 2010.
- David Elson and Kathleen McKeown. Automatic attribution of quoted speech in literary narrative. In *Proceedings of AAAI*, 2010.
- David Elson, Nicholas Dames, and Kathleen McKeown. Extracting social networks from literary fiction. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 138–147, Uppsala, Sweden, July 2010. Association for Computational Linguistics.
- Richard M. Emerson. Power-dependence relations. *American Sociological Review*, 27: 31–40, 1962.
- Bonnie Erickson, E Allan Lind, Bruce C Johnson, and William M O’Barr. Speech style and impression formation in a court setting: The effects of ”powerful” and ”powerless” speech. *J. Experimental Social Psychology*, 14(3):266 – 279, 1978.
- Kathleen Ferrara. Accommodation in therapy. In *Accommodation theory: Communication, context, and consequences*. Cambridge University Press, 1991.
- Eric Gilbert and Karrie Karahalios. Predicting tie strength with social media. In *CHI*, pages 211–220, 2009.
- Howard Giles. Communication accommodation theory. In *Engaging theories in interpersonal communication: multiple perspectives*. Sage Publications, 2008.

- Howard Giles, Justine Coupland, and Nikolas Coupland. Accommodation theory: Communication, context, and consequences. In *Accommodation theory: Communication, context, and consequences*. Cambridge University Press, 1991.
- Amy L. Gonzales, Jeffrey T. Hancock, and James W. Pennebaker. Language style matching as a predictor of social dynamics in small groups. *Communication Research*, 37(1):3–19, 2010.
- Stanford W. Gregory Jr. and Stephen Webster. A nonverbal signal in voices of interview partners effectively predicts communication accommodation and social status perceptions. *Journal of Personality and Social Psychology*, 70(6):1231, 1996.
- R. Guha, Ravi Kumar, Prabhakar Raghavan, and Andrew Tomkins. Propagation of trust and distrust. In *Proceedings of the 13th international conference on World Wide Web*, pages 403–412, 2004.
- Jerold L. Hale and Judee K. Burgoon. Models of reactions to changes in nonverbal immediacy. *Journal of Nonverbal Behavior*, 8(4):287–314, 1984.
- Jeffrey T. Hancock, Lauren E. Curry, Saurabh Goorha, and Michael Woodworth. On Lying and Being Lied To: A Linguistic Analysis of Deception in Computer-Mediated Communication. *Discourse Processes*, 45(1):1–23, 2008.
- Timothy Hawes. Computational analysis of the conversational dynamics of the United States Supreme Court. Master’s thesis, University of Maryland, 2009.
- Susan C. Herring. Gender and power in on-line communication. In Janet Holmes and Miriam Meyerhoff, editors, *The Handbook of Language and Gender*, pages 202–228. 2003.
- Susan C. Herring and John C. Paolillo. Gender and genre variation in weblogs. *Journal of Sociolinguistics*, 10(4):439–459, 2006.

- David Holmes. Authorship attribution. *Computers and the Humanities*, 28(2):87–106, Apr 1994.
- Courtenay Honeycutt and S.C Herring. Beyond microblogging: Conversation and collaboration via twitter. *HICSS*, 2009.
- Molly E. Ireland and James W. Pennebaker. Language style matching in writing: synchrony in essays, correspondence, and poetry. *J Personality and Social Psychology*, 99(3):549–571, 2010.
- Molly E. Ireland, Richard B. Slatcher, Paul W. Eastwick, Lauren E. Scissors, Eli J. Finkel, and James W. Pennebaker. Language style matching predicts relationship initiation and stability. *Psychological Science*, 22:39–44, 2011.
- Joseph Jaffé and Stanley Feldstein. *Rhythms of dialogue*. Academic Press, 1970.
- Akshay Java, Xiaodan Song, Tim Finin, and Belle Tseng. Why we twitter: understanding microblogging usage and communities. *WebKDD/SNA-KDD*, Aug 2007.
- Patrick Juola. *Authorship Attribution*. Now Publishers, 2008.
- Dan Jurafsky, Rajesh Ranganath, and Dan McFarland. Extracting social meaning: Identifying interactional style in spoken conversation. In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 638–646, Boulder, Colorado, June 2009. Association for Computational Linguistics.
- Justine Kao and Dan Jurafsky. A Computational Analysis of Style, Affect, and Imagery in Contemporary Poetry. In *Proceedings of the NAACL Workshop on Computational Linguistics for Literature*, 2012.

- Emre Kiciman. Language differences and metadata features on twitter. *Web N-gram Workshop*, Jan 2010.
- Philip Koehn. Statistical significance tests for machine translation evaluation. In *EMNLP*, pages 388–395, 2004.
- Moshe Koppel, Shlomo Argamon, and Anat Shimoni. Automatically categorizing written texts by author gender. *Literary and Linguistic Computing*, 2002.
- John P. Kotter. Power, dependence, and effective management. *Harvard Business Review*, pages 125–136, July 1977.
- Sarah Kozloff. *Overhearing Film Dialogue*. University of California Press, 2000.
- Jessica L. Lakin, Valerie E. Jefferis, Clara Michelle Cheng, and Tanya L. Chartrand. The chameleon effect as social glue: Evidence for the evolutionary significance of nonconscious mimicry. *Journal of Nonverbal Behavior*, 27:145–162, 2003.
- David Laniado, Riccardo Tasso, Yana Volkovich, and Andreas Kaltenbrunner. When the Wikipedians talk: Network and tree structure of Wikipedia discussion pages. In *ICWSM*, 2011.
- Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Governance in Social Media: A case study of the Wikipedia promotion process. In *ICWSM*, 2010a.
- Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. Predicting positive and negative links in online social networks. In *WWW*, pages 641–650, 2010b.
- Willem J M Levelt and Stephanie Kelter. Surface form and memory in question answering. *Cognitive Psychology*, 14(1):78 – 106, 1982.
- Adam Liptak. Clarence Thomas Keeps 5-Year Supreme Court Silence. *The New York Times*, February 2011.

- Yue Lu, Panayiotis Tsaparas, Alexandros Ntoulas, and Livia Polanyi. Exploiting social context for review quality prediction. *WWW*, pages 691–700, 2010.
- François Mairesse and Marilyn A Walker. Controlling User Perceptions of Linguistic Style: Trainable Generation of Personality Traits. *Computational Linguistics*, 2011.
- François Mairesse, Marilyn A. Walker, Matthias R. Mehl, and Roger K. Moore. Using linguistic cues for the automatic recognition of personality in conversation and text. *JAIR*, pages 457–500, 2007.
- Joseph Matarazzo and Arthur Wiens. *Interview: Research on Its Anatomy and Structure*. Chicago: Aldine, 1973.
- Andrew McCallum, Xuerui Wang, and Andr’es Corrada-Emmanuel. Topic and role discovery in social networks with experiments on Enron and academic email. *JAIR*, 30(1):249–272, 2007.
- Robert McKee. *Story: Substance, Structure, Style, and the Principles of Screenwriting*. Methuen, 1999.
- Rada Mihalcea and Carlo Strapparava. The lie detector: Explorations in the automatic recognition of deceptive language. In *Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 309–312, Suntec, Singapore, August 2009. Association for Computational Linguistics.
- Frederick Mosteller and David Wallace. Inference in an authorship problem. *Journal of the American Statistical Association*, 58(302):275–309, Jun 1963.
- Arjun Mukherjee and Bing Liu. Improving gender classification of blog authors. *EMNLP*, 2010.
- Randall Munroe. <http://xkcd.com/813/>, 2010.

- Laura L. Namy, Lynne C. Nygaard, and Denise Sauerteig. Gender Differences in Vocal Accommodation. *Journal of Language and Social Psychology*, 21(4):422, 2002.
- Michael Natale. Convergence of mean vocal intensity in dyadic communication as a function of social desirability. *J. Personality and Social Psych.*, 32(5):790–804, 1975.
- Kate G. Niederhoffer and James W. Pennebaker. Linguistic style matching in social interaction. *J. Lang. and Social Psych.*, 21(4):337–360, 2002.
- Jon Oberlander and Alastair J. Gill. Language with character: A stratified corpus comparison of individual differences in e-mail communication. *Discourse Processes*, 42(3):239–270, 2006.
- Myle Ott, Yejin Choi, Claire Cardie, and Jeffrey T. Hancock. Finding deceptive opinion spam by any stretch of the imagination. In *Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies*, pages 309–319, Portland, Oregon, USA, June 2011. Association for Computational Linguistics.
- Jahna Otterbacher and Libby Hemphill. Learning the lingo? Gender, prestige and linguistic adaptation in review communities. In *Proceedings of CSCW*, 2012.
- John C Paolillo. Language variation on Internet Relay Chat: A social network approach. *Journal of Sociolinguistics*, 5(2):180–213, 2001.
- James W. Pennebaker. *The secret life of pronouns: What our words say about us*. Bloomsbury Press, 2011.
- James W. Pennebaker, Roger J. Booth, and Martha E. Francis. Linguistic inquiry and word count (LIWC): A computerized text analysis program. <http://www.liwc.net/>, 2007.

- Cyma Van Petten and Marta Kutas. Influences of semantic and syntactic context on open- and closed-class words. *Memory and Cognition*, 1991.
- Martin Pickering and Simon Garrod. Toward a mechanistic psychology of dialogue. *Behavioral and Brain Sciences*, 27(02):169–190, 2004.
- Daniel Ramage, Susan Dumais, and Dan Liebling. Characterizing microblogs with topic models. In *International AAAI Conference on Weblogs and Social Media*, 2010.
- Rajesh Ranganath, Dan Jurafsky, and Dan McFarland. It’s not you, it’s me: Detecting flirting and its misperception in speed-dates. In *Proceedings of the 2009 Conference on Empirical Methods in Natural Language Processing*, pages 334–342, Singapore, August 2009. Association for Computational Linguistics.
- Paul Rayson, Andrew Wilson, and Geoffrey Leech. Grammatical word class variation within the British National Corpus Sampler. *Language and Computers*, 36:295–306(12), 2001.
- David Reitter, Frank Keller, and Johanna D Moore. Computational modelling of structural priming in dialogue. In *Proceedings of the Human Language Technology Conference of the NAACL, Companion Volume: Short Papers on XX*, pages 121–124, 2006.
- David Reitter, Frank Keller, and Johanna D Moore. A Computational Cognitive Model of Syntactic Priming. *Cognitive Science*, 2011.
- Alan Ritter, Colin Cherry, and Bill Dolan. Unsupervised modeling of Twitter conversations. In *Human Language Technologies: The 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 172–180, 2010.

- Ryan Rowe, Germán Creamer, Shlomo Hershkop, and Salvatore J Stolfo. Automated social hierarchy detection through email network analysis. In *Proceedings of the 9th WebKDD and 1st SNA-KDD 2007 workshop on Web mining and social network analysis*, pages 109–117, 2007.
- Ruchita Sarawgi, Kailash Gajulapalli, and Yejin Choi. Gender attribution: tracing stylistic evidence beyond topic and genre. In *Proceedings of the Fifteenth Conference on Computational Natural Language Learning*, pages 78–86, 2011.
- Antonin Scalia. Interview with Justice Antonin Scalia (transcript). C-SPAN (2009-06-19), 2009.
- Andrew J. Scholand, Yla R. Tausczik, and James W. Pennebaker. Social language network analysis. In *CSCW*, pages 23–26, 2010.
- Sidney J. Segalowitz and Korri Lane. Perceptual fluency and lexical access for function versus content words. *Behavioral and Brain Sciences*, 27(02):307–308, 2004.
- Morgan Sonderegger. Applications of graph theory to an English rhyming corpus. *Computer Speech & Language*, In Press, Corrected Proof, 2010.
- Svetlana Stenchikova and Amanda Stent. Measuring adaptation between dialogs. In *Proc. of SIGdial*, 2007.
- Richard L. Street and Howard Giles. Speech accommodation theory. In *Social cognition and communication*. Sage Publications, 1982.
- Fred L. Strodbeck and Richard D. Mann. Sex Role Differentiation in Jury Deliberations. *Sociometry*, 19(1):3–11, 1956.
- Devi Stuart-Fox and Adnan Moussalli. Selection for social signalling drives the evolution of chameleon colour change. *PLoS Biol*, 6(1):e25, 01 2008.

- Dario Taraborelli and Giovanni Luca Ciampaglia. Beyond notability. Collective deliberation on content inclusion in Wikipedia. *Second international workshop on quality in techno-social systems*, 2010.
- Yla R. Tausczik and James 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.
- Paul J Taylor and Sally Thomas. Linguistic style matching and negotiation outcome. *Negotiation and Conflict Management Research*, 1:263–281, 2008.
- Jitendra N. Thakerar, Howard Giles, and Jenny Cheshire. Psychological and linguistic parameters of speech accommodation theory. In C. Fraser and K.R. Scherer, editors, *Advances in the Social Psychology of Language*. Cambridge, 1982.
- Shane Thye. A status value theory of power in exchange relations. *American Sociological Review*, 65:407–432, June 2000.
- Fernanda B Viégas, Martin Wattenberg, Jesse Kriss, and Frank van Ham. Talk before you type: Coordination in Wikipedia. In *HICSS*, pages 78–78, 2007.
- Arthur Ward and Diane Litman. Dialog convergence and learning. In *Proceedings of the 2007 conference on Artificial Intelligence in Education: Building Technology Rich Learning Contexts That Work*, pages 262–269, Amsterdam, The Netherlands, The Netherlands, 2007. IOS Press. ISBN 978-1-58603-764-2.
- Stanley Wasserman and Katherine Faust. *Social network analysis: Methods and applications*. Cambridge University Press New York, 1994.
- Sheida White. Backchannels across cultures: A study of Americans and Japanese. *Language in Society*, 18(1):59–76, Mar 1989.

- Rebecca J. Wolfe and Kathleen L. McGinn. Perceived relative power and its influence on negotiations. *Group Decision and Negotiation*, 14:3–20, 2005.
- Fang Wu and Bernardo A. Huberman. Opinion formation under costly expression. *ACMTrans. Intelligent Systems and Technology*, 1(1):1–13, 2010.
- Patrick Ye and Timothy Baldwin. Verb sense disambiguation using selectional preferences extracted with a state-of-the-art semantic role labeler. In *Proceedings of the Australasian Language Technology Workshop 2006*, pages 139–148, Sydney, Australia, November 2006.
- G. Udny 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.