

Social Dynamics in Dialog:

A Case Study of the Enron Email Corpus Using Cultivate

A White Paper

Contents

Introduction	3
Previous Work	3
Case Study Methods and Results	4
Citations	6

Introduction

For sociolinguistic researchers studying the interplay of social context and conversational behavior, the digital interactions between members of large organizations or purely online communities are enticing datasets. Large and naturally occurring, these conversational corpuses offer social scientists the opportunity to conduct studies on a scale previously infeasible.

More recently, the advent of advanced computational techniques allowed for analyses on such corpuses to move beyond structural (network) features and into textual content. Researchers have found that social dynamics - most notably along dimensions of gender, gender environment, and power - manifest itself in dialog. Social context affects patterns of behavior that are not only detectable and quantifiable, but also distinct enough to allow accurate prediction of the governing social context from text alone.

We have assembled a collection of research papers which offer interesting results in the computational analysis of organizational behavior. While these papers all focus on variation in language use conditioned by social dynamics, the corpuses considered by each differ widely. The overarching goal of this case study is to reproduce these findings by using Cultivate's homegrown text classifiers and social structure analysis models on the gender-identified Enron email corpus.

In this white paper, we will first provide a brief overview of the papers and their findings. Then we will describe our methods for reproducing these findings using models we've built here at Cultivate, and we will describe the results.

Previous Work

Danescu-Niculescu-Mizil et al. (1) quantify the linguistic style of both discussions between Wikipedia editor and also transcripts of verbal exchanges before the U.S. Supreme Court. The authors measure levels of "language coordination," or the subconscious phenomenon between speakers within a power differential where one mimics the linguistic tendencies of those they are communicating with (in this case, admins and non-admins or Justices and lawyers). They identify two types of power - social status and social dependency. In social relationships defined by the former, non-admin users (low status) coordinate more towards admins (high status), and lawyers (low status) coordinate more towards Justices (high status). In social relationships defined by the latter, they find that lawyers (dependent) coordinate more towards unfavorable justices, and that admins-to-be (dependent) coordinate more towards all users (admin and non-admin) prior to a promotion, but less afterwards.

Choudhury et al. (2) evaluate affective expression by employees via an internal microblogging tool at a Fortune 500 company, showing changes in affect by time of day and organizational structure. Employees show a cyclic rhythm of affect throughout the day, peaking near noon and dropping significantly after-hours. Sales and marketing employees show higher affect than developers and researchers. Managers are also found to exhibit “affect accommodation,” or the tendency to adapt expressions of affect to the identity of the listener. By viewing microblogging content as a proxy for the previously inaccessible casual “watercooler conversation,” the authors succeed in identifying the effects of various endogenous and exogenous factors on employee emotion.

Prabhakaran et al. (3) provide the sociolinguistic research community a valuable augmentation of the Enron email dataset by identifying the gender of a large number of contacts within the corpus. Combining previous work in categorizing employee power dynamics, the authors are able to group conversations by both gender and power. They find that social dynamics evoke a number of conversational and textual cues - such as how female superiors use fewer overt displays of power than their male counterparts - then go even further to accurately predict power structure from dialogue alone.

Case Study Methods and Results

Using the Cultivate conversation analysis pipeline, we reproduced many of the findings above solely on the gender-identified Enron email corpus. Some of the required NLP methods, such as sentiment and style detection, were either available out-of-the-box or easily reproducible. A few others, including politeness and dialog act classifiers, were trained in-house to match or outperform the reported accuracies of the research models. Additionally, as part of the analysis, we’ve extracted latent network structure from group chats by fitting Network Hawkes models; i.e. extracting individual initiation and response tendencies, then clustering messages into likely conversations. Let’s go through the results!

First, we binned emails by hour of day and applied a sentiment analysis model, which ranks messages from low to high sentiment on a scale of 1 to 5. We found a similar cyclic rhythm as mentioned in (2) over a 24 hour period, with messages

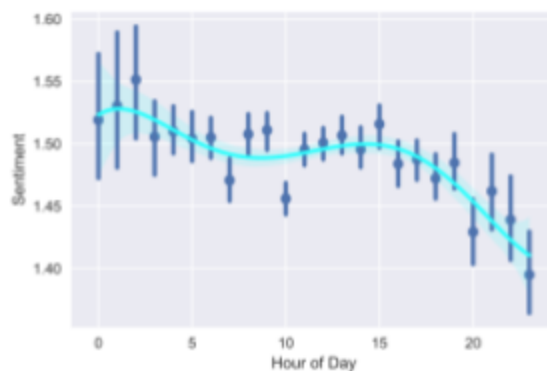


Figure 1. Sentiment by hour of day.

exhibiting fairly stable sentiment throughout the workday, but dropping off in the evening after work hours (see Figure 1).

Next, we used the linguistic style features in (1) to measure language coordination within the Enron dataset. Managers hold power over subordinates by virtue of their status in the workplace, and we saw that power differential reflected in coordination level. We found that subordinates are more likely to use certain linguistic elements while conversing with their manager, if the manager used that element in the immediately preceding message. The inverse effect (managers towards subordinates) is not as strong (see Figure 2).

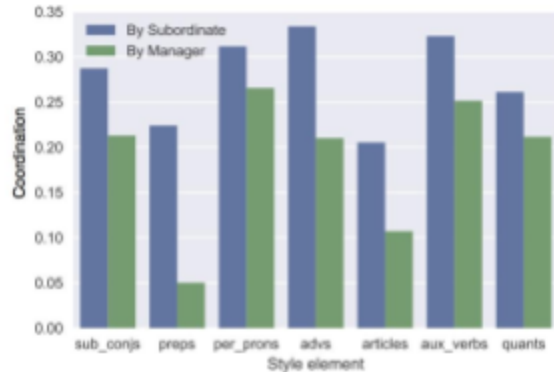


Figure 2. Subordinates coordinate more to managers than conversely.

Lastly, inspired by (3), we split contacts into four groups by gender and power (male/female, superiors/subordinates). Messages were classified into three broad dialog actions: request (either for information or action), inform (conveyance of information), and conventional (greetings, introductions, expressions of gratitude). We also tagged threads by gender environment, i.e. the dominant gender of thread participants. Results showed that subordinates contribute more information irrespective of gender environment, however the ratio of information imparted is skewed by gender environment (see Figure 3). In other words, the difference between superiors and subordinates in quantity of information imparted is much less in a male environment than a female environment. We also found that female superiors were the most polite in making requests, while male superiors were the least (see Figure 4).

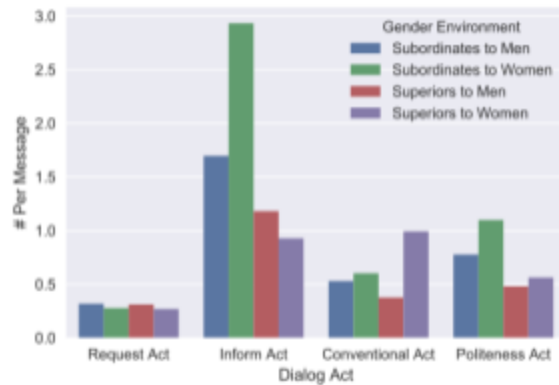


Figure 3. Messages classified by Dialog Act and Gender Environment.

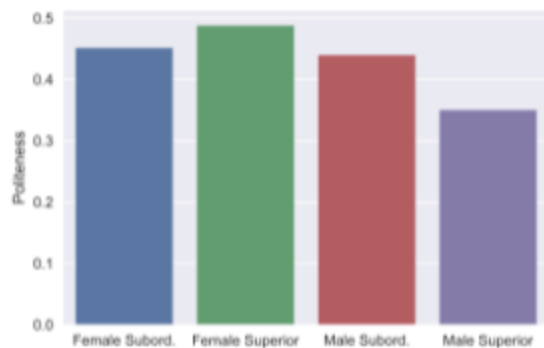


Figure 4. Politeness levels across gender environments.

Citations

- (1) Cristian Danescu-Niculescu-Mizil, Lillian Lee, Bo Pang: **“Echoes of power: Language effects and power differences in social interaction”**, 2011; [arXiv:1112.3670](https://arxiv.org/abs/1112.3670).
- (2) M. D. Choudhury and S. Counts: **“Understanding affect in the workplace via social media.”** In Proc. of the intl. conf. on Computer Supported Cooperative Work (CSCW), pages 303–316, 2013.
- (3) Vinodkumar Prabhakaran: **“Dialog Structure Through the Lens of Gender, Gender Environment, and Power”**, 2017, Journal for Dialogue & Discourse 8(2) (2017) 21-55; [arXiv:1706.03441](https://arxiv.org/abs/1706.03441). DOI: [10.5087/dad.2017.202](https://doi.org/10.5087/dad.2017.202).