



CS 6474/CS 4803 Social Computing: Analyzing Language II

Munmun De Choudhury


munmund@gatech.edu

Week 7 | February 25, 2021


Gender and Power: How
Gender and Gender
Environment Affect
Manifestations of Power

Summary

- Interplay between gender, gender environment of online conversations and power
- Contributions:
 - Automatic gender assignment of 87% of the Enron corpus – US Social Security Administration list of names matching the approx. age range of Enron employees; first names gathered from email headers
 - Test the sociolinguistic hypotheses: face-saving use of language, and to the use of language to strengthen social relations
 - Gender-based features boosts the accuracy of predicting the direction of power between pairs of email interactants



Overt display of power



Conventional Counts

Would these results hold in modern enterprise social media contexts?

When Social Networks Cross Boundaries: A Case Study of Workplace Use of Facebook and LinkedIn

Meredith M. Skeels

Biomedical & Health Informatics, University of Washington
Seattle, WA, USA
mskeels@u.washington.edu

Jonathan Grudin

Microsoft Research
Redmond, WA, USA
jgrudin@microsoft.com

ABSTRACT

The use of social networking software by professionals is increasing dramatically. How it is used, whether it enhances or reduces productivity, and how enterprise-friendly design and use might evolve are open questions. We examine attitudes and behaviors in a large, technologically-savvy organization through a broad survey and thirty focused interviews. We find extensive social and work uses, with complex patterns that differ with software system and networker age. Tensions arise when use spans social groups and the organization's firewall. Although use is predominantly to support weak ties whose contribution to productivity can be difficult to prove, we anticipate rapid uptake of social networking technology by organizations.

Categories and Subject Descriptors

H.5.3. Group and Organization Interfaces.

to corporate use. One-third of the employees in the enterprise we studied were in the Facebook company network. We found an equal number of employees with LinkedIn accounts. Professional-oriented LinkedIn had quadrupled in size to over 25 million members in one year [12]. What if anything are all these professional users doing with social networking software?

Enterprise adoption of social networking software is far easier, and preventing it more difficult, than was true for earlier technologies. This raises questions. Do these sites enhance productivity? Can utility for enterprises be increased? What new issues will arise for these new user populations?

In early 2008, we conducted this research in Microsoft, then an organization of 88,000. Although not a typical enterprise, it is typical of the early adopters of email and IM that foreshadowed subsequent wider use. We briefly review social networking

A Longitudinal Study of Facebook, LinkedIn, & Twitter Use

Anne Archambault
Microsoft Corporation
Redmond, Washington USA
annea@microsoft.com

Jonathan Grudin
Microsoft Research
Redmond, Washington USA
jgrudin@microsoft.com

ABSTRACT

We conducted four annual comprehensive surveys of social networking at Microsoft between 2008 and 2011. We are interested in how these sites are used and whether they are considered to be useful for organizational communication and information-gathering. Our study is longitudinal and based on random sampling. Between 2008 and 2011, social networking went from being a niche activity to being very widely and heavily used. Growth in use and acceptance was not uniform, with differences based on gender, age and level (individual contributor vs. manager). Behaviors and concerns changed, with some showing signs of leveling off.

Author Keywords

Social networking; Facebook; LinkedIn; Twitter; Enterprise

messaging, and employee blogging were first used mainly by students and consumers to support informal interaction. Managers, who focus more on formal communication channels, often viewed them as potential distractions [4]. A new communication channel initially disrupts existing channels and creates management challenges until usage conventions and a new collaboration ecosystem emerges.

Email was not embraced by many large organizations until the late 1990s. Instant messaging was not generally considered a productivity tool in the early 2000s. Slowly, employees familiar with these technologies found ways to use them to work more effectively. Organizational acceptance was aided by new features that managers appreciated, such as email attachments and integration with calendaring.

Microblogging Inside and Outside the Workplace

Kate Ehrlich, N. Sadat Shami

IBM TJ Watson Research Center and Center for Social Software
1 Rogers St., Cambridge, MA 02142, USA
{katee, sadat} @ us.ibm.com

Abstract

Microblogging has recently generated a lot of research interest. Yet very little is known about how corporate employees use microblogging tools. This study examined microblogging in the workplace by conducting a content analysis comparing posts from individuals who were using an internal proprietary tool and Twitter simultaneously. In both settings, posts that provided information or were directed to others were more common than posts on status. Within these categories, it was more frequent to provide information externally than internally but more common to ask questions either through broadcast or directed posts internally than externally. Qualitative interviews explored users' motivations regarding microblogging behavior. The paper concludes with a discussion of the implications of microblogging for business use.

Grudin 2007; Skeels and Grudin 2009), especially with respect to issues of privacy and confidentiality.

This study sought to understand the use and value of microblogging in the workplace by analyzing over 5000 microblog posts from a group of employees. These employees were unique in that they used both an internal proprietary tool to post internally and Twitter to post externally. By comparing internal and external posts from the same users we explore how posts directed at a workplace only audience might differ from posts that are directed to a broader audience of non-work as well as work colleagues. Interviews with a majority of these employees allowed us to learn more about their reasons for posting internally or externally, and the value they got from posting and reading microblogs.

Table 1

Prospects for the Study of Enterprise Social Media Use in Organizations

Metaphors of Enterprise Social Media			
Processes	Leaky Pipe	Echo Chamber	Social Lubricant
Social Capital	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Easy to “keep up” with what others are doing without significant social investment. • Broad knowledge helps build bridges across nonredundant groups. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Awareness that others see what/whom you know could stop you from contributing so as not to undermine brokerage position. • Potential loss of power from making private rolodexes public 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Immediate feedback from similar others strengthens existing communities. • Helps to establish common ground that makes interaction and sense of belonging easier. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Self-reinforcing groups may balkanize and splinter into non-redundant communities. • Groupthink could arise from exposure only to similar others 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Insights into what others are doing and who they know help create conversational fodder that makes it easy to initiate new connections and maintain established connections. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Peripheral awareness of others may create illusion that a real social connection exists when it does not. • Too much social information can disrupt work and distract from work-related communication.
Boundary Work	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Ability to cross more knowledge boundaries due to visibility into what people are doing in other groups, departments, or locations. • Ability to see more connections between people and forge alliances. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Understanding of people in different parts of the organization, but doing similar tasks, can increase sense of relationships and belonging. • Promote similarity and accessibility in global teams, across cultures. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Ease of communication creates a low stakes environment to reach out to people not within same social group. • Blurring boundaries between private and work communication showcase personal similarities that can be touch points for work communication.
	<p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • More generic communication due to awareness that people outside a trusted or known community are watching. • Loss of proprietary information in a particular group. 	<p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Strengthen boundaries between groups making communication, interaction and identification more difficult. • Create a “speaker’s corner” in which people only from one side of boundary interact and listen to each other 	<p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Context collapse makes it difficult to know which “self” to present in what situation. • Highlights differences in communication style across cultures, which can make people more reticent to reach out across boundaries
Attention Allocation	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Individuals begin to attend to information, knowledge, and communication from others who they would not normally talk with. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Many information inputs means cognitive overload and individuals allocate attention only to specific areas of the organization, or discontinue use of ESM altogether due to overload. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Because of public nature of communication to a known community, people provide more accurate and honest information. • Information from trusted others increases attention to ideas communicated by others. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Individuals may believe that information they are attending to is representative of entire organization. • Construction of sub-optimal attention allocation strategies. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Due to threaded and temporally ordered nature of conversation, people can focus their attention in ways that allows them to enter conversations more easily at meaningful times. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • People interject in conversations not intended for them. • Too many social-related signals can scatter one’s attention and increases absentmindedness.
Social Analytics	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Because communication is visible and available, managers can use these digital traces to understand the organization’s informal information economy • Create strategic opportunities for connecting people who are not yet connected <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Increased ability for surveillance and possibility of control. • Knowledge that management is watching may compel people to refrain from communicating on the platform. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Better understand who are the various communities within the organization, even if those communities are not tied to formal organizations (e.g. departments or divisions). <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Mistaken understanding of what communities are or who key players in them might be because analytics do not sample communication that occurs offline. 	<p><i>Advantages</i></p> <ul style="list-style-type: none"> • Recommendations of connections provides excuse for people to get to know one another • Recommendations for documents that one might read can provide conversation-starter material with documents’ creators. <p><i>Disadvantages</i></p> <ul style="list-style-type: none"> • Encourages strategic self-presentation or offline interactions to avoid being traced, tracked, and quantified, which reduces likelihood people will use the tool to make new connections.

As people in corporations increasingly adopt platforms like Facebook and Twitter, how do you expect these findings to generalize/change?

How can the sociolinguistic findings on gender, gender environment, and power be useful for social computing research?

Class Exercise I

What kind of design considerations could incorporate the sociolinguistic findings on gender, gender environment, and power so that they aid people in the maintenance of social relations as Prabhakaran et al found? You can describe using any social computing platform of choice.

How do the preexisting social/corporate structures and biases at a large corporation impact the data, especially given that less than 35% of the unique discourse participants were women?

The style of communication depends on the context or environment

Man is to Computer Programmer as Woman is to Homemaker? Debiasing Word Embeddings

Tolga Bolukbasi¹, Kai-Wei Chang², James Zou², Venkatesh Saligrama^{1,2}, Adam Kalai²

¹Boston University, 8 Saint Mary's Street, Boston, MA

²Microsoft Research New England, 1 Memorial Drive, Cambridge, MA

tolgab@bu.edu, kw@kwchang.net, jamesyzou@gmail.com, srv@bu.edu, adam.kalai@microsoft.com

Abstract

The blind application of machine learning runs the risk of amplifying biases present in data. Such a danger is facing us with *word embedding*, a popular framework to represent text data as vectors which has been used in many machine learning and natural language processing tasks. We show that even word embeddings trained on Google News articles exhibit female/male gender stereotypes to a disturbing extent. This raises concerns because their widespread use, as we describe, often tends to amplify these biases. Geometrically, gender bias is first shown to be captured by a direction in the word embedding. Second, gender neutral words are shown to be linearly separable from gender definition words in the word embedding. Using these properties, we provide a methodology for modifying an embedding to remove gender stereotypes, such as the association between the words *receptionist* and *female*, while maintaining desired associations such as between the words *queen*

Semantics derived automatically from language corpora necessarily contain human biases

Aylin Caliskan¹, Joanna J. Bryson^{1,2}, and Arvind Narayanan¹

¹Princeton University

²University of Bath

*Address correspondence to aylinc@princeton.edu, bryson@conjugateprior.org, arvindn@cs.princeton.edu.

ABSTRACT

Artificial intelligence and machine learning are in a period of astounding growth. However, there are concerns that these technologies may be used, either with or without intention, to perpetuate the prejudice and unfairness that unfortunately characterizes many human institutions. **Here we show for the first time that human-like semantic biases result from the application of standard machine learning to ordinary language—the same sort of language humans are exposed to every day.** We replicate a spectrum of standard human biases as exposed by the Implicit Association Test and other well-known psychological studies. We replicate these using a widely used, purely statistical machine-learning model—namely, the GloVe word embedding—trained on a corpus of text from the Web. Our results indicate that language itself contains recoverable and accurate imprints of our historic biases, whether these are morally neutral as towards insects or flowers, problematic as towards race or gender, or even simply veridical, reflecting the *status quo* for the distribution of gender with respect to careers or first names. These regularities are captured by machine learning along with the rest of semantics. In addition to our empirical findings concerning language, we also contribute new methods for evaluating bias in text, the Word Embedding Association Test (WEAT) and the Word Embedding Factual Association Test (WEFAT). Our results have implications not only for AI and machine learning, but also for the fields of psychology, sociology, and human ethics, since they raise the possibility that mere exposure to everyday language can account for the biases we replicate here.



Social media is distorting the representation of women in Africa. Here's what can be done about it



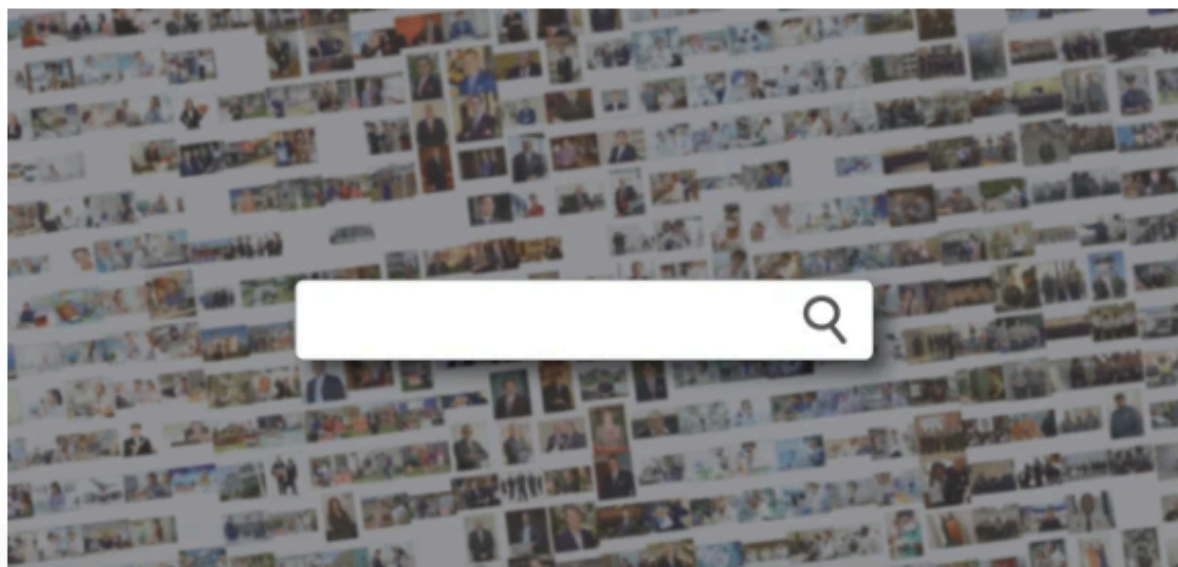
DECEMBER 17, 2018



Gender and Jobs in Online Image Searches

Men are overrepresented in online image search results across a majority of jobs examined; women appear lower than men in such search results for many jobs

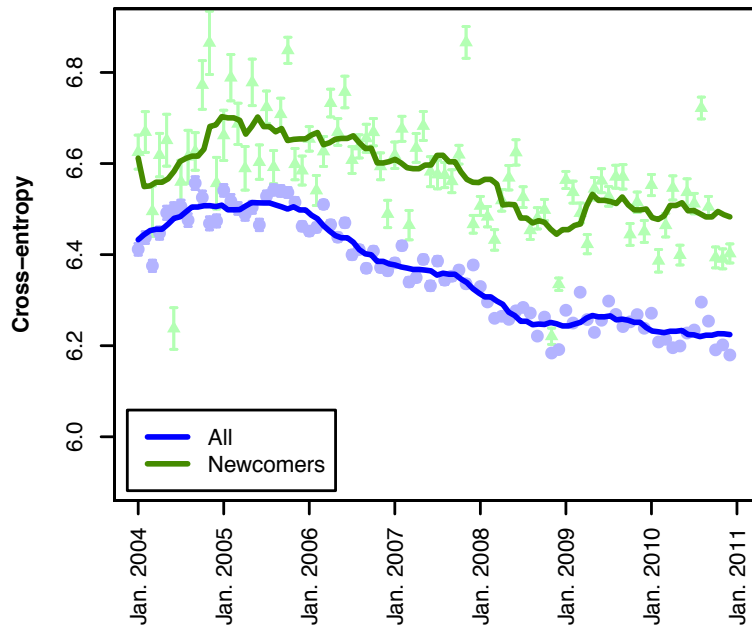
BY ONYI LAM, BRIAN BRODERICK, STEFAN WOJCIK AND ADAM HUGHES




No Country for Old Members:
User lifecycle and linguistic
change in online communities

Summary


- The paper proposes a framework for tracking linguistic change as it happens in a community, to understand how specific users react to the community's evolving norms
- Results show a two-stage lifecycle of linguistic change in communities (RateBeer and BeerAdvocate):
 - a linguistically innovative learning phase in which users adopt the language of the community
 - a conservative phase in which users stop changing and the evolving community norms pass them by




(a) BeerAdvocate



(b) RateBeer



(a) BeerAdvocate



(b) RateBeer

Community Acceptance and Feedback

Why is this important from the perspective of the social computing field –class exercise and brainstorming examples later...

Class Exercise II

Danescu-Niculescu-Mizil et al. say that “[their] framework can be used to detect, early in a user’s career, how long she will stay active in the community”

Describe two scenarios where this knowledge will be beneficial and two where it will not be useful/could be harmful. Who are these stakeholders who will be interested in this knowledge?

BUILDING MEMBER ATTACHMENT IN ONLINE COMMUNITIES: APPLYING THEORIES OF GROUP IDENTITY AND INTERPERSONAL BONDS¹

Yuqing Ren

Carlson School of Management, University of Minnesota, Minneapolis, MN 55455 U.S.A. {chingren@umn.edu}

F. Maxwell Harper, Sara Drenner, Loren Terveen

Department of Computer Science, University of Minnesota, Minneapolis, MN 55455 U.S.A.

{harper@cs.umn.edu} {mosch@cs.umn.edu} {terveen@cs.umn.edu}

Sara Kiesler

Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA 15213 U.S.A. {kiesler@cs.cmu.edu}

John Riedl

Department of Computer Science, University of Minnesota, Minneapolis, MN 55455 U.S.A. {riedl@cs.umn.edu}

Robert E. Kraut

Human-Computer Interaction Institute, Carnegie Mellon University, Pittsburgh, PA 15213 U.S.A. {robert.kraut@cmu.edu}

Predicting Continued Participation in Newsgroups

Elisabeth Joyce

Edinboro University of Pennsylvania

Robert E. Kraut

Human-Computer Interaction Institute
Carnegie Mellon University

Turnover in online communities is very high, with most people who initially post a message to an online community never contributing again. In this paper, we test whether the responses that newcomers receive to their first posts influence the extent to which they continue to participate. The data come from initial posts made by 2,777 newcomers to six public newsgroups. We coded the content and valence of the initial post and its first response, if it received one, to see if these factors influenced newcomers' likelihood of posting again. Approximately 61% of newcomers received a reply to their initial post, and those who got a reply were 12% more likely to post to the community again; their probability of posting again increased from 44% to 56%. They were more likely to receive a response if they asked a question or wrote a longer post. Surprisingly,

Feed Me: Motivating Newcomer Contribution in Social Network Sites

Moira Burke

Human-Computer Interaction Institute
Carnegie Mellon University
5000 Forbes Ave., Pittsburgh, PA 15213
moira@cmu.edu

Cameron Marlow and Thomas Lento

Facebook
156 University Ave.
Palo Alto, CA 94301
{cameron, lento}@facebook.com

ABSTRACT

Social networking sites (SNS) are only as good as the content their users share. Therefore, designers of SNS seek to improve the overall user experience by encouraging members to contribute more content. However, user motivations for contribution in SNS are not well understood. This is particularly true for newcomers, who may not recognize the value of contribution. Using server log data from approximately 140,000 newcomers in Facebook, we predict long-term sharing based on the experiences the newcomers have in their first two weeks. We test four mechanisms: social learning, singling out, feedback, and distribution.

is a *content feed*, which publishes stories about a user or set of users and makes the stories available to others. Such feeds may cause users to increase their rate of content contribution, either by increasing user awareness of product features and the socially acceptable means of using them, encouraging users to contribute content to attract the attention of their peers, or a combination of these effects.

This paper examines the relationship between initial user behavior and content production in a social network environment. Using a set of approximately 140,000 Facebook users who joined in March 2008, we examine the newcomers' initial content contribution and their friend networks to assess the effects of friends' behavior, feedback,

Mental Health Support and its Relationship to Linguistic Accommodation in Online Communities

Eva Sharma
Georgia Tech
Atlanta, GA USA
evasharma@gatech.edu

Munmun De Choudhury
Georgia Tech
Atlanta, GA USA
munmund@gatech.edu

ABSTRACT

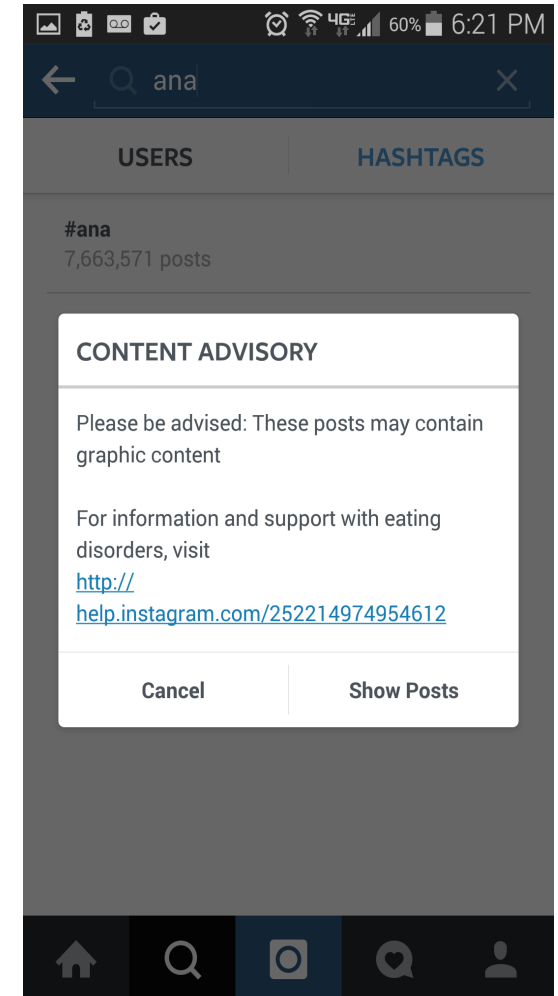
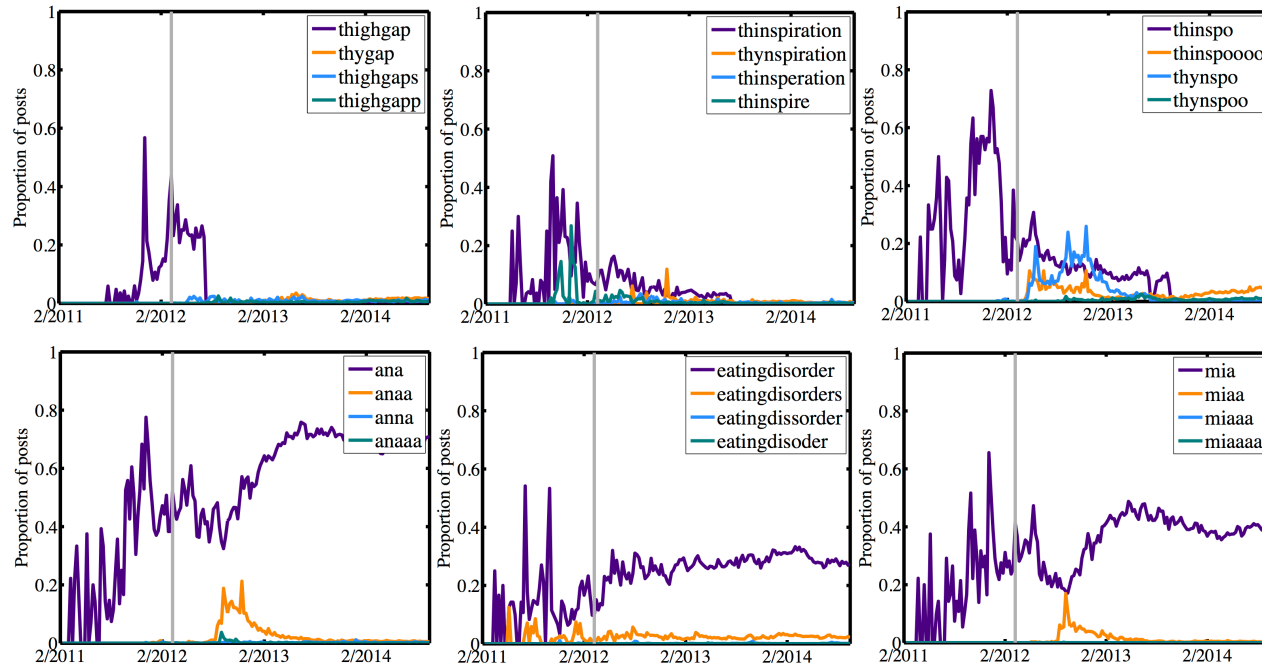
Many online communities cater to the critical and unmet needs of individuals challenged with mental illnesses. Generally, communities engender characteristic linguistic practices, known as norms. Conformance to these norms, or linguistic accommodation, encourages social approval and acceptance. This paper investigates whether linguistic accommodation impacts a specific social feedback: the support received by an individual in an online mental health community. We first quantitatively derive two measures for each post in these communities: 1) the linguistic accommodation it exhibits, and 2) the level of support it receives. Thereafter, we build a statistical framework to examine the relationship between these measures. Although the extent to which accommodation is associated with support varies, we find a positive link between the two, consistent across 55 Reddit communities serving various psychological needs. We discuss how our work surfaces

by providing a ‘buffer’ against the potentially adverse effects of stressful or difficult situations [20, 57]. However, outside of therapeutic contexts, vulnerable individuals often have limited ability to access adequate social support [50, 51, 78].

Online mental health communities (OMHCs), in recent years, have emerged as prominent resources for mental health support [84]. In fact, support derived from these communities has been found to causally improve mental wellbeing like reduced likelihood of suicidal thoughts [31]. Such support can range from *emotional support* (ES) to *informational support* (IS), often taking the form of empathy, acknowledgment, advice, or situational appraisal around diverse issues like mental illness, crisis, addiction, and abuse [17, 30, 31]. Moreover, due to the high quality of support provided by these OMHCs, they are also considered as a “safe haven”: they enable individuals to express disinhibiting emotions, engage in self-disclosures

Will the two-phase lifecycle (linguistic innovation learning and conservative phases) hold for/generalize to other online communities?

Chancellor, S., Pater, J. A., Clear, T., Gilbert, E., & De Choudhury, M. (2016, February). *# thyghgapp: Instagram content moderation and lexical variation in pro-eating disorder communities.* In Proc. CSCW.



#anorexia, #anarexia, #anarexyia: Characterizing Online Community Practices with Orthographic Variation

Ian Stewart, Stevie Chancellor, Munmun De Choudhury, and Jacob Eisenstein
Georgia Institute of Technology
Atlanta, GA

{istewart6, schancellor3, munmun.choudhury, jacob.eisenstein}@gatech.edu

Abstract—Distinctive linguistic practices help communities build solidarity and differentiate themselves from outsiders. In an online community, one such practice is variation in *orthography*, which includes spelling, punctuation, and capitalization. Using a dataset of over two million Instagram posts, we investigate orthographic variation in a community that shares pro-eating disorder (pro-ED) content. We find that not only does orthographic variation grow more frequent over time, it also

members, then what characterizes the members who accept and advance these changes?

The social meaning of language change in online communities can be better understood by linking language change to community membership dynamics, i.e., the progression of individual community members from new to

What about anonymous communities (like 4chan /b/) where it may be less important whether someone is a newcomer?

What about communities catering to specific needs of people (e.g., mental health)?