Mind Your P's and Q's: When Politeness Helps and Hurts in Online Communities

Moira Burke
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15217 USA
moira@cmu.edu

Robert Kraut
Human-Computer Interaction Institute
Carnegie Mellon University
Pittsburgh, PA 15217 USA
robert.kraut@cmu.edu

CHI 2008, April 5 – April 10, 2008, Florence, Italy
ACM 978-1-60558-012-8/08/04.

Abstract
Little is known about the impact of politeness in online communities. This project combines deductive and inductive approaches to automatically model linguistic politeness in online discussion groups and determine the impact of politeness on desired outcomes, such as getting people to reply to one another. We find that politeness triples reply rates in some technical groups, but rudeness is more effective in others. The model can be applied as a "politeness checker" to encourage people to write in ways likely to garner a response from specific communities.

Keywords
Linguistic politeness, computer-mediated communication, community responsiveness

ACM Classification Keywords
H.5.3 [Information Interfaces]: Group and Organization Interfaces - Collaborative computing, Web-based interaction, Computer-supported cooperative work

Introduction
Though our mothers advised us to mind our p's and q's, little is known about the effect of politeness in computer-mediated communication. This is especially true for online communities, in which people attempt to start conversations and make requests of strangers. Does polite conflict resolution lead a Wikipedia editor to be promoted to admin status? Do polite responses to newcomers in health support groups cause those newcomers to help others in the future? Does it get you killed in World of Warcraft?

To answer these questions, we need ways to measure politeness in online communities. The current project combines two approaches to train a machine learning algorithm to automatically model polite language: (1) A deductive approach based on Brown and Levinson's linguistic politeness theory [2], and (2) an inductive, bag-of-words method that identifies words and short phrases perceived as polite. The present paper describes the extension of theory developed thirty years ago to a coding manual for online discussion groups, and a survey of perceived politeness, in which current newsgroup messages are rated, providing a gold standard for a bottom-up model of polite language. Using reply count data for these messages, we determined the impact of politeness on community responsiveness, and found large differences based on community type: Politeness tripled reply counts in some technical groups, while rudeness was more effective in garnering replies in some political issue groups.

Linguistic Politeness Theory

Face and Linguistic Politeness Strategies

Linguistic politeness theory begins with Goffman's theory of "face" [7], the presentation of an identity with positive social value, and it is important to have that identity validated by others. However, in the presence of others, we are subject to face-threatening actions, such as impositions and criticisms. So, we put in face-work when communicating in order to help maintain each other's identities. Grice's maxims for efficient conversation describe the most direct forms of speech (e.g. "Take out the trash.") [8]. Yet these forms are often lengthened in everyday conversation (e.g. "Would you please take out the trash?"), and Brown and Levinson propose that this inefficiency is an attempt to save one another's face [2]. By being indirect, we "implicate" some degree of politeness, which the hearer recognizes while still understanding the underlying meaning of the utterance[9].

Based on observations of language in three cultures, Brown and Levinson describe a typology of linguistic politeness strategies. In this paper, we focus on two categories of strategies: "negative politeness" in which the speaker attempts to minimize the imposition on the listener (e.g. "If you have chance, would you close the window?"), and "positive politeness" indicating a social connection between the speaker and listener (e.g. "Let's close the window."). Specific strategies from each category are listed below.

Negative politeness strategies:
N1. Be conventionally indirect
N2. Question, hedge
N3. Be pessimistic
N4. Minimize the imposition
N5. Give deference
N6. Apologize
N7. Impersonalize the speaker and hearer
N8. State the face threatening action as a general rule
N9. Nominalize
N10. Go on record as incurring a debt

Positive politeness strategies:
P1. Notice, attend to the hearer's needs
P2. Exaggerate interest, approval, sympathy
P3. Intensify interest to the hearer
P4. Use in-group identity markers
P5. Seek agreement
P6. Avoid disagreement
P7. Presuppose/raise/assert common ground
P8. Joke
P9. Assert/presuppose knowledge of hearer's concerns
P10. Offer, promise
P11. Be optimistic
P12. Include both speaker and hearer in activity
P13. Give or ask for reason
P14. Assume or assert reciprocity
P15. Give gifts to the hearer

There are two main criticisms of Brown and Levinson's model. First, the strategies are ambiguous, partially overlapping, and fall at many different levels of communication, from syntactic (e.g. question-form) to pragmatic (e.g. joking) [14]. The lowest-level strategies may be relatively easy to detect automatically, while the higher-level ones will be difficult even for human coders. Second, the focus is on the speaker's perception of politeness, rather than the recipient's [14]. Yet speakers and writers often overestimate their ability to convey subtler cues, such as sarcasm [11]. Given this overestimation and myriad cultural norms regarding politeness, it is possible that intended politeness is not always received. Therefore, the current project matches ratings of perceived politeness with intended politeness.

Politeness Research in CMC

Politeness research in computer-mediated communication generally falls into two camps: One set applies small subsets of Brown and Levinson's typology to medium-sized corpora, while the other camp applies all or most of the typology to very small datasets.[1] The small subsets of Brown and Levinson's typology commonly applied are those terms that are easily countable, such as "please," "thank you," "would," and hedges [3][13]. Brennan and O'Haeri counted hedges and questions in instant messaging, and suggested that the belief that people sound less polite in CMC can be attributed to production costs: It takes more time to type hedges and indirect requests in fast-paced CMC, and so people use balder, shorter forms [1]. Yet adding a question mark takes little extra effort, so question forms were as common in instant messaging as in face-to-face communication. These studies have successfully applied small, easily countable portions of Brown and Levinson's model to computer-mediated communication, but have not delved into a more comprehensive application of the model, as the current project proposes.

A few studies have applied Brown and Levinson's model in its entirety to very small datasets from specific domains. Carlo and Yoo compared transcripts of 14 face-to-face transactions between reference librarians and students to 15 reference sessions via online chat [5]. They found significantly more negative and fewer positive politeness strategies online than in face-to-face transactions. Simmons coded ten weeks of messages from an online bulletin board on censorship and described face-threatening actions, most of which were threats to negative face [12]. He suggests that over time people will show more positive face-saving strategies online, as people adjust to this "faceless" medium, though this is contrary to Carlo and Yoo's findings. Duthler compared the politeness strategies used in email to voicemail when students had to make low- and high-imposition requests of a fictitious professor, and found differences in strategies within email between low- and high-imposition requests, but no differences in voicemail messages, suggesting that email is more tailorable, though extraneous phrases were correlated with decreased perceived politeness [6]. In general, both camps of politeness research in CMC are descriptive, but few connect politeness strategies with desired outcomes, such as getting the receiver to reply, or to contribute to the community in the future. The current project will determine the impact of specific politeness strategies on these desired outcomes.

Method

To build a model of linguistic politeness, we harvested a set of 576 messages posted to 12 discussion groups from 2004 to 2006. The groups cover a wide variety of topics, including diabetes, depression, multiple sclerosis, atheism, economics, life extension, C programming, math,

[1] Additional studies of politeness in CMC have focused on gender, culture, or other domains, such as computer tutoring, but these are outside the scope of this project.

electronics design, piloting, quilting, and general discussion by people over fifty. Each message was the first in its thread, and thus an attempt to start conversation, rather than a reply to an ongoing conversation. The number of replies to each message was counted, and usernames and signatures were replaced with same-gender pseudonyms.

For each of these messages we need two pieces of information: how polite readers perceive it to be, and which linguistic politeness strategies it contains. To measure perceived politeness, we surveyed 194 readers, described below. We have also developed a coding manual extending Brown and Levinson's politeness strategies to computer-mediated communication also described briefly below, and are in the process of coding each message for those strategies.

Measuring Perceived Politeness

To measure perceived politeness, we recruited 225 participants for a thirty-minute web-based survey with a random raffle for one of five $80 gift certificates. The survey was advertised in online classified ads across the U.S., and participants were required to answer 4 of 5 randomly selected grammar questions from the Test of English as a Foreign Language (TOEFL) correctly to proceed. Each participant read 48 randomly ordered messages counterbalanced from the 12 discussion groups and rated each message on a seven-point scale from very rude to very polite. Thirty-one participants were excluded for finishing the survey in less than 15 minutes, including two who selected "4" on the politeness scale for nearly every message, leaving 194 participants. Overall, an average of 14 politeness ratings per message were gathered, with good inter-rater reliability (Cronbach's alpha = 0.93 and mean correlation between any two judges is 0.41). To account for individual biases, each individual's score for a message was standardized by subtracting her mean score across all messages and divided by her standard deviation. Each message's perceived politeness score was then the mean of the standardized scores from each participant.

Identifying Linguistic Politeness Strategies

To identify specific politeness strategies, such as hedging or seeking agreement, a coding manual has been created that extends Brown and Levinson's 25 positive and negative strategies to online discussion groups, adding keywords and examples from current groups. Each code will be applied independently, allowing for overlapping codes within a message (e.g. joking, apologizing, and minimizing imposition). Examples for two strategies are included below (bold added to highlight strategy):

N6. Apologize

*"So, if you could, please send whatever healing energies you can. . . . This is the scariest thing I've ever been through . . . Thanks in advance and **apologies for the imposition.**"*

*"**normally wouldn't ask but** a few days ago, I made the announcement that PIF had broken the 3000 pipe sent 'barrier', I thought this was kinda cool, but not one response"*

P2. Exaggerate interest, approval, or sympathy

*"I'm still glad you had an **AWESOME** time. That concert was so RARE and you guys got something a lot of fans will never get. **CONGRATS!!**"*

*"**Congrats on the Mac**. . . It will help with your quilting. Really. I promise**! ;-P**"*

Results

We found significant differences both in the perceived politeness in the 12 newsgroups and the impact of politeness on community responsiveness. Table 1 shows the results of a linear regression on politeness, with dummy variables for the groups, and controlling for the message length (Mean=654.5 characters, SD=603.3,

Table 1. Linear regression on politeness in the 12 groups. Coefficients represent the number of standard deviations (SDs) from the mean. Thus, messages to the math group were 1.39 SDs more polite than messages posted to the omitted group, atheism.

	Coefficient	Std. Err.
Intercept	-1.23***	.20
Message length in characters (ln)	0.06*	.03
Groups		
Economics	-0.08***	0.12
Life extension	0.69***	0.12
Diabetes	0.99***	0.12
Depression	0.67***	0.12
Multiple Sclerosis	1.20***	0.12
C programming	1.06***	0.12
Math	1.39***	0.12
Electronics design	1.17***	0.12
Aviation	0.78***	0.12
Quilting	1.32***	0.12
Over-50 chat	0.82***	0.12
N=576 messages	*** p < .001	** p < .01 * p <.05

normalized by adding 1 and taking the natural log in the regression). Overall, the issue groups were perceived as least polite, and math and quilting most polite.

To determine the impact of politeness on community responsiveness, we performed a negative binomial regression on the number of replies messages in those groups received. Negative binomial regressions are appropriate for non-negative count variables with overdispersion, as is commonly the case in online discussion groups. A binary dummy variable "Is Polite" has the value of 1 if the message's standardized politeness score was positive (and thus more polite than average), and 0 otherwise. Table 2 shows that there is an interaction between politeness and group: Controlling for overall differences in reply rates between groups, politeness triples reply counts in math and programming groups, while rudeness triples replies in the atheism group.

The number of replies can be calculated with the following equation:

$$\textit{Number of replies} = \beta_0 + \beta_{IsPolite}(\textit{Is Polite}) + \beta_{Group}(\textit{Group}) + \beta_{IsPoliteXGroup}\ (\textit{Is Polite x Group}) + \beta_{Length}(\textit{Message Length})$$

Thus, a rude message of mean length posted to the omitted group, atheism, would receive: 2.90 + 0.16 = 3.06 replies. A polite message posted to that group would receive roughly one-third the number of replies: 2.90 – 2.15 + 0.16 = 0.91 replies. By comparison, rude messages posted to the math group received 0.20 replies, while polite messages received 0.68, more than three times as

Table 2. Negative binomial regression on the number of replies messages in each group received, controlling for message length and group. Atheism is the omitted group.

	Coefficient	Std. Err.
Intercept	2.90***	0.40
Message length in characters (ln)	0.16**	0.06
Is Polite	-2.15***	0.64
Groups		
Economics	-1.58***	0.22
Life extension	-2.88***	0.28
Diabetes	-0.75**	0.28
Depression	-1.57***	0.24
Multiple Sclerosis	-1.65***	0.35
C programming	-1.82***	0.33
Math	-2.86***	0.44
Electronics design	-1.10***	0.31
Aviation	-0.66**	0.26
Quilting	-1.13**	0.39
Over-50 chat	-0.91***	0.26
Interactions		
Economics X Is Polite	1.37 n.s.	0.86
Life extension X Is Polite	2.48***	0.73
Diabetes X Is Polite	2.07**	0.71
Depression X Is Polite	1.93**	0.72
Multiple Sclerosis X Is Polite	2.10**	0.74
C programming X Is Polite	2.32***	0.73
Math X Is Polite	2.63***	0.78
Electronics design X Is Polite	1.66*	0.73
Aviation X Is Polite	1.60*	0.71
Quilting X Is Polite	2.01**	0.76
Over-50 chat X Is Polite	2.13**	0.71
N=576 messages	*** p < .001	** p < .01 * p <.05

many responses. Politeness nearly tripled responses in the life extension group, as well, from 0.18 to 0.51. Though these effect sizes sound small, previous research has shown that receiving even a single response has dramatic outcomes on an individual's future behavior, including increased likelihood of posting again [4] and a greater participation answering others' questions in the future[10].

Conclusion

The overall goal of this work is to determine how politeness affects the experiences that people have in online communities. To understand that, we are building a model of linguistic politeness driven both by theory and bottom-up perceptions of politeness. Using a survey to obtain politeness scores for a set of messages on diverse topics, we found large differences in perceived politeness between communities and in the effectiveness of politeness on reply counts.

Previous politeness research has relied upon human codes of small sets of data; this project includes a machine leaner that can be applied to much larger corpora, for greater generalizability and the design of automatic interventions, such as a "politeness checker" that suggests linguistic strategies to newcomers before they post their first messages. The machine learner can also be applied to other kinds of messages, such as replies, to determine if people who receive polite replies in their early group interactions go on to contribute more to the group in the future (such as replying to others). The learner can be applied to other domains, such as Wikipedia or SourceForge, to determine if politeness in production communities leads to greater or higher-quality products. Automatically detecting linguistic politeness in online communities will increase our understanding of how strangers make successful requests and become integrated into communities through conversation.

Acknowledgments

The authors thank Sue Fussell, Alla Zats, and Neel Shah. This work was supported by NSF IIS-0325049 and an NSF Graduate Research Fellowship.

References

[1] Brennan, S., & Ohaeri, J. (1999). Why do electronic conversations seem less polite? The costs and benefits of hedging. *Proceedings of International Joint Conference on Work Activities, Coordination, and Collaboration (pp. 227-235).* San Francisco, CA: ACM.

[2] Brown, P., & Levinson, S. (1987). *Politeness: Some Universals in Language Usage*. Cambridge: Cambridge University Press.

[3] Bunz, U., & Campbell, S. (2004). Politeness accommodation in electronic mail. *Communication Research Reports*, 21(1), 11-25.

[4] Burke, M., Joyce, E., Kim, T., Anand, V., and Kraut, R. (2007). Introductions and requests: Rhetorical strategies that elicit response in online communities. *Proc. Communities & Technologies*.

[5] Carlo, J., & Yoo, Y. (in press). "How may I help you?" Politeness in computer-mediated and face-to-face reference transactions. *Information and Organization*.

[6] Duthler, K. W. (2006). The politeness of requests made via email and voicemail: Support for the hyperpersonal model. *Journal of Computer-Mediated Communication*, 11, Article 6. http://jcmc.indiana.edu/vol11/issue2/duthler.html

[7] Goffman, E. (1959) *The Presentation of Self in Everyday Life*. Garden City, NY: Doubleday Anchor.

[8] Grice, H. (1957). Meaning. *Philosophical Review*, 67, 377-388.

[9] Holtgraves, T. (2002). Language as Social Action: Social Psychology and Language Use. Mahwah, NJ: Erlbaum.

[10] Joyce, E. and Kraut, R. Predicting continued participation in newsgroups. *Journal of Computer Mediated Communication* 11,3(2006).

[11] Kruger, J., Eppley, N., Parker, J., & Ng, Z-W. (2005). Egocentrism over email: Can we communicate as well as we think? *Journal of Personality and Social Psychology,* 89, 925-936.

[12] Simmons, T. (1994*). Politeness Theory in Computer-Mediated Communication: Face-Threatening Acts in a "Faceless" Medium*. Unpublished masters thesis, Asthon University, Birmingham, England. Retrieved April 23, 2007 from http://eric.ed.gov/ERICWebPortal/contentdelivery/servlet/ERICServlet?accno=ED381005

[13] Spertus, E. (1997). Smokey: Automatic Recognition of Hostile Messages. *Proc. AAI 1997* Boston: AAAI Press.

[14] Watts, R. (2003). Modeling linguistic politeness (II): Brown and Levinson and their critics. In *Politeness*, (pp. 85-116). Cambridge: Cambridge University Press.