Etiquette and Politeness in Human-Human and Human-Machine Interactions:

A Summary of Work at SIFT

Christopher A. Miller
Smart Information Flow Technologies (SIFT)
Minneapolis, MN U.S.A.
cmiller@sift.info
+1 612 716 4015

ABSTRACT

This paper provides an overview, with multiple references, to work on "etiquette" in human-machine interaction, and the development of a computational model to characterize, quantify and simulate it, as performed by individuals at Smart Information Flow Technologies (SIFT). Our work began with the realization that operators personify and interact with complex automation according to human-like behavioral conventions even if that automation is in no way explicitly embodied or personified. This has led us to interest in characterizing and quantifying aspects of culture-specific human-human etiquette in a fashion usable in computer systems. In particular, we have developed a computational model of politeness in interaction based on the sociolinguistic work of Brown and Levinson. We have recently demonstrated integration of our model into a game-like simulation, providing characters the ability to both recognize and generate culture-specific polite or impolite behaviors, and have provided both validation data and demonstration of substantial savings in software engineering costs with our model. We are now beginning work using our model to predict (again, culture-specific) human compliance with "directives" using different levels and types of politeness and delivered either by other humans or by a machine.

Author Keywords

Etiquette, Politeness, Human-Computer Interaction, Culture.

INTRODUCTION

For the past several years we've been interested in what we are calling "etiquette" in human interactions—first in human-machine interaction and more recently in human-human, cross-cultural interactions. "Etiquette" is a term we chose consciously in part to be provocative (whoever heard of a machine caring about etiquette!), and also because of its association with protocols that carry substantial social meaning even though they may be unwritten and largely implicit.

ETIQUETTE IN HUMAN-MACHINE INTERACTIONS

We came upon this topic, perhaps, through the back door beginning with an interest in human-machine interactions [5]. For many years we have been involved in the creation of complex and sophisticated automation to support difficult and high criticality real-world tasks such as piloting aircraft [6], operating oil refineries [2] and caring for the sick and elderly [7]. In most of these domains, highly-trained human operators are not interested in, and at least believe that they will not benefit from, embodied interfaces with explicit humanlike characteristics. Nevertheless, we began to realize that such operators were approaching our complex and sophisticated automation with behaviors and expectations that were similar to those they would use with other humans, even in spite of our best attempts to avoid such personification. Eventually we found the work of Clifford Nass [14] which provided both a theoretical framework and substantial experimental data to justify this claim.

We began to believe that such responses were inevitable in any even moderately sophisticated form of automation, and therefore that we should take such phenomena into account in our designs. This does not necessarily lead us to embodied interfaces, but rather to what we called taking an "etiquette perspective" on design—realizing that humans have a tendency to personify and therefore designing around that tendency either encouraging or discouraging aspects of it as they may facilitate trust, safety, accuracy of performance, efficiency, etc.

In one early case, the Cockpit Information Manager we designed for the Rotorcraft Pilot's Associate (a smart, task-based display manager for the pilots of a dual-crew attack/scout helicopter) [6], this philosophy led us away from what had been a typical and exclusive reliance on plan recognition capabilities, into a more interactive and conversational approach to intent “declaration” or specification--even though that approach might arguably produce higher workload. This was precisely because intent declaration was what was expected in the “etiquette” of this cockpit by these pilots. The result of this approach was improved user acceptance and arguably, improved overall human plus machine performance [6]. In another, subsequent study [13], we demonstrated that poor “etiquette” in the recommendations that a decision aid provides can have as big an impact on user trust, perceived workload and, most important, overall human + machine task performance as the difference between 60% and 80% reliability in the accuracy of the aid’s recommendations. That is, users who got advice from the decision aid that was only accurate 60% of the time, but got it in a polite, non-interruptive and non-nagging fashion, did as well on an objective performance-based diagnosis task as those who got 80% accurate advice, but got it from a “rude”, nagging decision aid.

A COMPUTATIONAL MODEL OF POLITENESS

Brown and Levinson’s Qualitative Politeness Model

An emphasis on “etiquette” has almost inevitably led us to an emphasis on cross-cultural representations and interactions—since etiquette inevitably, perhaps even by definition, differs from culture to culture and subculture to subculture. Approximately four years ago, while searching for more quantitative and explicit models of aspects of etiquette we came upon the work of Brown and Levinsion [1], two sociolinguistics who had extensively studied and developed a universal (although qualitative) model of politeness in human-human interactions. Brown and Levinson were interested in explaining the “inefficiency” of the politeness behaviors which appear in all human languages and cultures. When I say “please pass the salt,” the use of “please” is not necessary for truthful, relevant or clear expression of my wish and is, in fact, not required to express my overt intent. Brown and Levinson collected a large database of instances of politeness in communication across three major cultural/linguistic groups (English, Tamil and Tzetzal) and, from this data, developed a qualitative model which both identified cross-cultural commonalities in politeness behaviors and proposed a culturally universal model of how, when and why politeness is used.

Their explanation for politeness behaviors stems from the fact that humans are intentional agents with the potential to have their will, intentions, and sense of self-worth or -regard (that is, their “face” [3]) threatened. Virtually all interactions between social agents are to some degree Face Threatening Acts (FTAs). My simple act of speaking to you, regardless of the content of my words, places a demand on your attention that threatens your ability to autonomously direct it wherever you want. This, then, is the reason for the “please” in my request for salt: If I simply state my desire (e.g., “Give me the salt”) I would be ambiguous about whether I have the power or right to compel you to give me salt and you might well take offense. “Please” (as for all politeness behaviors) is thus a “redressive” strategy which mitigates the threat. Furthermore, the expectation that such a strategy be used is an example of etiquette that enables interpretations. The etiquette is the “rule” that entitles us to conclude that those who use “please” are striving to be seen as polite; those who do not are not striving to be polite for various reasons (perhaps they don’t believe they need to be, perhaps their notions about politeness are different, perhaps they are just rude).

Our interpretation of Brown and Levinson’s qualitative model (see Figure 1) declares that an interaction between two individuals will be perceived as balanced or “nominally polite” if the degree of face threat in the interaction is balanced by the value of the set of politeness behaviors (aka “redressive actions”) used in the interaction. If more politeness is used than there was face threat in the interaction, then the interaction will be seen as “over polite”; if less politeness was used than there was threat then the interaction will be seen as “under polite” or rude.

Face threat itself, in the Brown and Levinson model, is a function of the observer’s perception of three additional parameters:

- P(H,S) is the relative power that the hearer (H) has over the speaker (S). Power is an asymmetric relationship. The higher the power of the hearer, the greater will be the face threat all things being equal, and thus the more redress will be required to balance out that threat. Thus, “give me the salt” might be appropriate for me to use to a low powered individual in our culture, such as a young grade school student, but very inappropriate for me to use to a high-powered individual, such as the

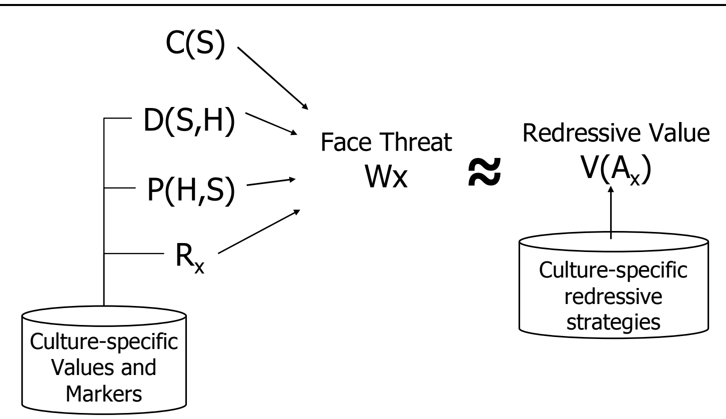

Figure 1. Our interpretation of Brown and Levinson's [1] politeness model.

CEO of a major corporation or the President of the United States.

- D(S,H) is the social distance between the speaker (S) and the hearer (H). Social distance is roughly the inverse of familiarity and is a symmetric relationship. The greater the social distance between speaker and hearer, and therefore the more redress required. Thus, "give me the salt" might be appropriate to an old friend, but very inappropriate to a complete stranger.
- R is the ranked imposition of the raw act itself. Some acts are simply more threatening than others. For example, asking for a loan of $500 is more threatening than asking for a loan of $.50. Thus, I can ask a complete stranger "can you tell me the time?" without much in the way of politeness behaviors, but if I wanted to ask for a larger favor, such as a ride to the airport, I might need to be much more polite: "I'm sorry, sir, but I'm in real trouble, I'd very much appreciate it if you could possibly give me a quick ride to the airport."

With regard to polite, redressive strategies, Brown and Levinson go on to identify some 40 general types of politeness behavior that they have observed across multiple cultures in their corpus. A few of these are illustrated in the last example given above:

- *Apology*—"I'm sorry…" Explicitly acknowledges the face threat and shows that I am contrite for it
- *Give deference*—"… sir…" Using an honorific explicitly builds up the face of the Hearer.
- *Give Reasons*—"… but I'm in real trouble…" Accounts for the face threat as stemming from other sources than my explicit intentions
- *Incur Indebtedness*—"… I'd very much appreciate it…" Acknowledges debt incurred by the FTA
- *Be pessimistic*—"…if…possibly…" Minimize face threat by leaving the compliance decision with H
- *Minimize imposition*—"… a quick ride…" Minimize face threat by implying that the magnitude of the imposition is small.

Note that the general dynamics of this model, and the specific concepts and terms that it uses, are intended to be culturally universal. That is, all cultures are believed to determine whether an utterance is polite or impolite in context based on whether their perception of the face threat present in the interaction is balanced by the amount of redress. Furthermore, all cultures are also presumed to reckon face threat as a function of power difference, social distance and degree of raw imposition. And finally, all cultures have been observed, in Brown and Levinson's work, to use redressive behaviors that fall into the categories described above. That said, what counts as a face threat, or what counts as power or social distance, and what counts as a specific instance of a general type of redressive behavior, as well as the value of all of these parameters, all differ from culture to culture. This is so obvious as to be almost missable in language: thus I say "sir" in English, while I say "saheb" in Pashto—two different sound patterns each of which is an instance of the general category "honorific" (and each of which may have a different weight in their respective cultures). Gestures behave similarly: taking off my hat is a sign of deference in many Western cultures, but would not be recognized as such in Iraq, where taking off one's sunglasses has roughly the same effect and weight.

Benefits from a Quantitative Politeness Model

Our work over the past few years has been seeking to develop a quantitative and computational model based on the qualitative model of Brown and Levinson. We have found methods for representing and quantifying the various aspects of the model described above: we have developed scales for representing power difference, social distance and raw imposition, as well as methods for identifying and scoring politeness behaviors (both verbal and non-verbal) along with an algorithm for combining them to assess their overall value in context. In addition, we have expanded the overall model in a few ways. For example, we have added a term representing what the observer knows about the "character" of the speaker: C(S). Detailed descriptions of our representation and its application in a variety of settings may be found in [8, 9, 10]. We have now completed several partial validation exercises involving this representation and algorithm, and can claim that it shows promising accuracy at least for American cultural interpretations [11].

There are several core benefits to be gained from a computational representation of politeness, especially one like Brown and Levinson's which is abstracted away from, yet is instantiatable by, culture-specific knowledge. As illustrated in Figure 2, we have now demonstrated the ability for the same core algorithm to both recognize politeness behaviors directed at it in a game-like setting, and to select politeness behaviors to be used in generating utterances and behaviors directed at others which are in keeping with its overall goals. Representing verbal versus nonverbal politeness behaviors is no challenge for our algorithm if they can first be recognized in the game or simulation setting in which the algorithm operates; they are both simply instances of redressive behaviors and can be scored and combined similarly. The character's perceptions and reactions are dictated by our politeness algorithm which operates over a culture-specific knowledge base (as illustrated in Figure 2). While the development of such knowledge bases is still a non-trivial amount of work—and the character's perceptions and reactions will only be as extensive as the knowledge represented therein, we will discuss ways in which our model and its algorithmic implementation streamlines knowledge capture and representation below. We have also demonstrated the ability for our core algorithm to be populated with culture-specific knowledge bases containing culture-specific values

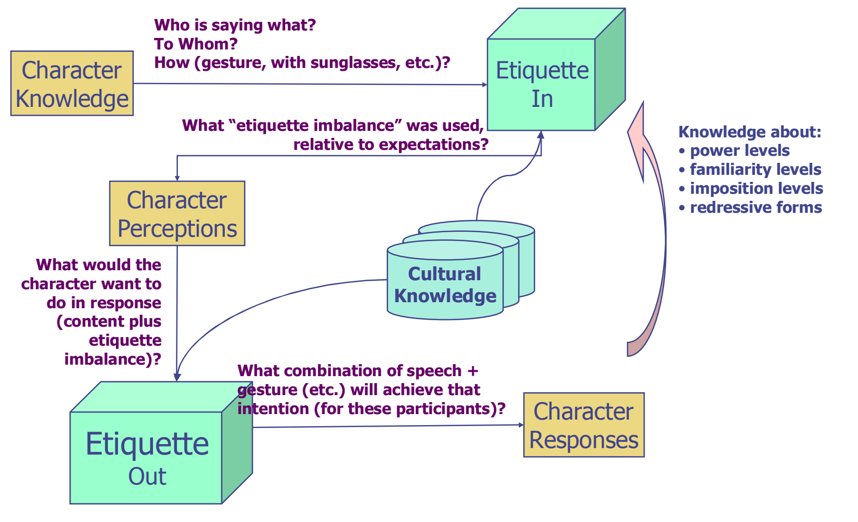

Figure 2. Integration of our etiquette algorithm in a character simulation environment.

for power relationships, social distance relationships, imposition scores as well as a culture – specific lexicon of politeness behaviors and their values. Such "cultural modules", once built (an important caveat, of course), enable us to change the cultural sensitivities (if not the look-and-feel) of a simulated character from, say, that of an Iraqi imam to an American private with the "flick" of a software switch. Since our algorithm operates identically over different sets of knowledge about these attributes, the character's sensitivities can be changed with ease.

For example, we have implemented the sunglasses vs. hat removal knowledge described above in two different knowledge bases corresponding to Iraqi politeness knowledge and American politeness knowledge. The knowledge that removing your hat is a sign of deference with a given redressive value is included in the American knowledge base, but absent from the Iraqi one, while the knowledge about leaving sunglasses on is present in the Iraqi knowledge base, but absent in the American one. Given these alternate sets of knowledge, it is trivial for us to show how a given action (say, uttering a greeting with hat on) will be perceived by an American character (somewhat rude) vs. an Iraqi (neutral—since there is no knowledge that hats "should" be removed to show deference) by simply running our algorithm with a different knowledge base of etiquette perceptions and values. In principle, this approach can be extended to the individual level—we could readily create a knowledge base for an individual who knows both sunglass and hat removal redressive behaviors: say, an American who has learned about this particular aspect of Iraqi etiquette.

Initial funding for this work has come from sources (primarily DARPA and the US Army) interested in computational and algorithmic representations of human cultural behavior for the purpose of computer-based training games and/or simulations. Our claim was that an algorithmic approach to recognizing and/or generating politeness behaviors would dramatically reduce the software engineering effort required to enable animated characters to exhibit such behaviors. This reduction in effort could then be used to either expand the range of behaviors that characters could exhibit (thus making them more realistic and believable), or to reduce the costs of developing such games, or both. This is clearly an advantage over traditional methods of enabling software simulations to exhibit culture-specific behaviors--particularly, hand-scripting chains of behavioral rules—and in fact, we have demonstrated [10] geometric increases in the speed with which we can add "Perception Scores" (PSs—how one observer perceives the politeness of one specific communication uttered by a specific speaker-hearer pair) to a simulation. We encoded our first set of 42 PSs at the rate of 2.33 PSs/hour, but the next set were acquired at 19.89 PSs/hour, and the final set of more than 2000 PSs were acquired at the rate of 48.96/hour.

ETIQUETTE AND DIRECTIVE COMPLIANCE

More recently, however, we have begun to turn our attention back to our original interests in the role of etiquette in human performance and in human-machine interactions. We are using our computational model of politeness as a guide to researching the effects of politeness on "directive compliance behavior." A "directive" is any utterance whose force is to direct the hearer to perform an action. Thus, a directive may be a request, instruction, command, or beseeching, depending on the etiquette with which it is delivered. We are interested, in short and for example, in whether saying "please" makes a difference on whether or not the hearer complies with the directive given, and if so, what kind of difference. For example, it may well be that the use of greater politeness in a directive results in greater overall probability that the hearer will comply with the directive, and in greater trust and liking for the directive giver, but it may also be the case that reaction times will be longer.

Furthermore, we are interested in the relationship between cultural factors as identified by authors such as Hofstede [4] and Nisbett [12] on the relationship between politeness perception and directive compliance behavior. For example, Hofstede identifies a cultural dimension he calls "Power Difference Index" (PDI) which refers to the degree to which members of a culture see large differences in power between individuals to be natural and acceptable versus abnormal or intolerable. We hypothesize that members of cultures with high PDI will be more tolerant of power differences in interactions, and thus will expect more

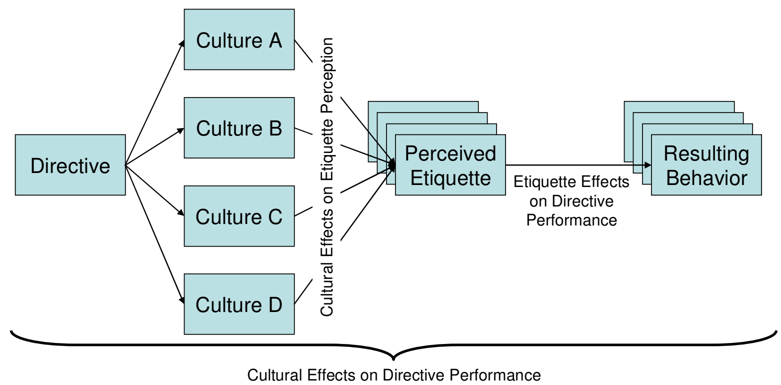

Figure 3. The CECAEDA chain by which culture-based perceptions of etiquette may affect compliance decisions and yield culture-specific behaviors.

potent redressive behaviors from a low-power speaker to a high-power hearer than will members from a high PDI culture observing the same interaction.

We have postulated a chain from perception through decision-making to behavior by which politeness may affect compliance behaviors and by which cultural factors may affect both. This chain is illustrated in Figure 3 and is known as our CECAEDA model, for "Computable Effects of Cultural Attributes and Etiquette on Directive Adherence". We are currently using this CECAEDA model to inform the design of two testbeds and a series of experiments to explore the effects of etiquette variations on directive compliance.

The first testbed, PAMMI ("Park Asset Monitoring and Management Interface"), will be used in work sponsored by the US Air Force to explore the effects of etiquette variations in directives given by other humans on compliance behavior from a multi-cultural population of experimental participants. PAMMI places the operator in the role of a dispatch officer monitoring the progress of a firefighting effort in a national park and relaying information about the location, path, progress, and role of assets to various individuals "in the field" who ask for it via a chat interface—albeit with varying degrees of politeness. We can control aspects of the power relationship, social distance, and even appearance and communication modality used by the information requesters, as well as the degree and type of politeness used in their requests, and we can assess whether subjects comply with the request, with what speed and accuracy, as well as their memory for the question and events surrounding it, their perceived trust or affect for the requestor, their perceived workload at the time, etc.

The second testbed is known as CHEWBACCA ("Computer Human Etiquette to Weight Benchmarking Assets and Capabilities to Corresponding Allocations"). CHEWBACCA has been designed to explore the effects of etiquette variations in directives from a decision aiding system on compliance behavior. In this testbed, which simulates the same fire fighting scenario as PAMMI, operators play a more active dispatcher role. Here, they must decide how to allocate airborne vehicles from a flight (each of which has different capabilities) to a series of active "targets" (each of which has different needs). For example, one target may be a group of trapped campers, thus their need is for a vehicle which is capable of transporting them. A vehicle which is only equipped with sensors will not do. The number and variety of vehicles and targets can be combined to make this task quite challenging for the operator, and a decision aiding system provides recommendations (which may or may not be correct or optimal) which may come to the operator in textual or audible form with varying degrees of politeness. Again, in this testbed we are interested in whether the politeness variations in the way in which the aid delivers its recommendations will make a difference in compliance behavior, reaction time, perceived workload, memory, trust, and affect.

Work is underway on experiments that use both testbeds, and we anticipate completing at least some of these experiments with their associated data analysis in time for the workshop.

ACKNOWLEDGMENTS

Work described in this document has been supported by contracts from DARPA (W31P4Q-04-C-R221), the U.S. Army (Contracts W31P4Q-04-C-R221 and DAAD-19-01-2-0009) and the U.S. Air Force (contract FA8650-06-C-6635). Contributors to this work include Lewis Johnson, Hannes Viljalmsson and Ning Wang of the University of Southern California, and Peggy Wu, Harry Funk, Marc Chapman and Peggy Wilson of SIFT participated in elements of this work.

REFERENCES

1. Brown, P. and Levinson, S. *Politeness: Some Universals in Language Usage,* Cambridge University Press, 1987.
2. Cochran, T., Miller, C. and Bullemer, P. Abnormal Situation Management in Petrochemical Plants: Can a Pilot's Associate Crack Crude? In *Proc. NAECON* 1996.
3. Goffman, E. *Interaction Ritual: Essays on Face to Face Behavior.* Garden City; 1967.

4. G. Hofstede, *Cultures and Consequences, 2nd Ed.* Sage Publications, 2001.
5. Miller, C. Rules of Etiquette, or How a Mannerly AUI should Comport Itself to Gain Social Acceptance and be Perceived as Gracious and Well-Behaved in Polite Society. In *Notes AAAI Spring Symp. on Adaptive User Interfaces*. 2000.
6. Miller, C. and Hannen, M. The Rotorcraft Pilot's Associate: Design and Evaluation of an Intelligent User Interface for a Cockpit Information Manager. *Knowledge Based Systems, 12* (1999), 443-456.
7. Miller, C., Wu, P., Krichbaum, K. and Kiff, L. Automated Elder Home Care: Long Term Adaptive Aiding and Support We Can Live With. In *Notes AAAI-Spring Symp. on Interaction between Humans and Autonomous Systems over Extended Operatio.,* 2004.
8. Miller, C., Chapman, M., Wu, P. and Johson, L. The "Etiquette Quotient": An Approach to Believable Social Interaction Behaviors. In *Proc. BRIMS* 2005.
9. Miller, C., Wu, P., Funk, H., Wilson, P, and Johnson, L. A Computational Approach to Etiquette and Politeness: Initial Test Cases. In *Proc. BRIMS* 2006.
10. Miller, C., Wu, P., Funk, H., Johnson, L. and Viljalmsson, H. A Computational Approach to Etiquette and Politeness: An "Etiquette Engine™" for Cultural Interaction Training. In *Proc. BRIMS* 2007.
11. Miller, C., Wu, P. and Funk, H. A Computational Approach to Etiquette and Politeness: Validation Experiments. In *Proc. 1st Int. Con. Computational Cultural Dynamics*. 2007.
12. Nisbett, R. *The geography of thought: How Asians and Westerners think differently…and why.* Free Press, 2003.
13. Parasuraman, R. and Miller, C. Trust and Etiquette in High-Criticality Automated Systems. *Comm. ACM. 47*, 2004, 51-55.
14. Reeves, B. and Nass, C. *The Media Equation.* Cambridge University Press, Cambridge, UK. 1996.