

Social Dialogue is Serious Business

Timothy Bickmore

MIT Media Lab

20 Ames St., Room E15-320

Cambridge, MA 02139

+1 617 253 7368

bickmore@media.mit.edu

ABSTRACT

The ability for social agents to establish personal relationships with their users is relevant not only in entertainment applications, but is important in most task-oriented applications, and is absolutely crucial in certain domains such as education, sales and psychotherapy. Following human behavior as a model, the primary mechanism that should be used to establish these instrumental relationships is social dialogue in the context of simulated face-to-face conversation. Past and present work to develop such "Relational Agents" is presented, along with suggestions for future research.

Keywords

Embodied conversational agent, trust, social interface, natural language.

INTRODUCTION

Over the last several years there has been a growing interest in social interfaces which leverage users' natural inclinations to respond psychosocially to computer generated stimuli as if they were produced by another person in the context of a face-to-face interaction. Examples include the classic work of Reeves & Nass [28], work in embodied conversational agents (ECAs) [8], and a series of additional studies and systems published at CHI and elsewhere (e.g., [25]). However, compared to typical human-human social interaction, these studies and systems have analyzed artifacts that recognize and display only the most minimal of social cues. In the commercial arena, people appear willing to engage artifacts such as Tamagotchis and Furbies in even more sophisticated and encompassing social interactions. This trend seems to indicate a human readiness, or even need, to engage computational artifacts in deeper and more substantive social interactions. Unfortunately, given the state of the art, current artifacts inevitably let users down. The systems are typically designed to interact with users for a very brief period of time, produce a very limited range of social cues and behaviors, ignore most or all of the user's social cues, and—worst of all—forget everything about the user they might have learned during their initial interaction.

In contrast, we humans use a wide range of behaviors and communicative modalities to build dynamic, on-going relationships with each other. We use the context of face-

to-face conversation to engage in small talk to build rapport and common ground, intimacy through self-disclosure, credibility through the use of expert's jargon, social networks through gossip, and "face" through politeness. We are able to engage not only in the Gricean maximally-informative task-oriented frame of interaction typically assumed in HCI [15], but are able to engage in play, joking, telling stories, teasing, and countless other frames of interaction. These frames are important, because they are the ones typically used to build and maintain relationships. However, they break many of the fundamental assumptions built into traditional task-oriented interfaces, such as the Gricean assumptions of relevance, non-redundancy, and truthfulness in communication (think about messages like "Wonderful weather we're having!"). We also use a myriad of nonverbal cues to display our attitudes towards each other (e.g., "immediacy" behaviors [22]) and to signal when talk switches from frames of social interaction to task interaction and back [16]. Our relationships persist over time, from interaction to interaction, and yet they are also dynamic in nature. We continually re-negotiate the status of our relationships, verbally and nonverbally proposing changes in status and reacting to similar proposals by our partner.

Why should we bother to build such capabilities into our social agents? In the remainder of this paper I will discuss the importance of personal relationships in human-human interactions, not only in purely social contexts, but in explicitly task-oriented contexts as well, and why it is important for our agents to be able to participate in these. I will then argue that social dialogue in the context of face-to-face interaction is the primary mechanism that people use to build and maintain their relationships, and how social agents can use these same techniques to move relationships with users in a desired direction. Finally, I will review previous and on-going work in building 'relational agents' which have some of these capabilities, and conclude with some observations about future directions for this work.

THE IMPORTANCE OF PERSONAL RELATIONSHIPS

Before discussing the instrumentality of personal relationships in arenas in which the more personal, non-task-oriented, aspects of relationship *do* play an important role, it is instructive to first look at the role of relationships

in arenas in which the more personal aspects of relationships are downplayed.

One example of such an arena is the world of corporate bureaucracy. Even in this area, the development of a network of interpersonal relationships has been found to be critical to a general manager's ability to implement his or her agenda, and the quality of these relationships has been found to be a key determinant of managerial effectiveness. In other studies, subordinates reporting good relationships with superiors have been found to be better performers, assume more responsibility and contribute more to their units than those reporting poor relationships [13].

In the study of service interactions, researchers differentiate between *service relationships*, in which a customer expects to interact again in the future with the same service provider (and vice versa), *pseudorelationships*, in which a customer expects to interact again in the future with the same firm (but not the same person), and *service encounters*, in which there are no such expectations of future interactions. In a series of surveys involving 1,200 subjects, Gutek, et al, found that subjects who are in service relationships reported more trust in and knowledge of their service providers, more interest in continuing the interaction, and more willingness to refer the provider to others, than customers in either pseudorelationships or service encounters [17]. The results also indicate that a service relationship with a particular human service provider is significantly more effective at engendering trust, commitment and referrals than attempts to establish brand or firm loyalty.

In addition to these more general work-related issues, there are some application domains in which the formation of a quality personal relationship is seen not as an added benefit, but as a prerequisite to a successful outcome. Examples include certain types of sales, education, and psychotherapy domains.

Within sales, the customer-salesperson relationship is important to successful sales outcomes, especially in major purchase situations such as real estate where there is both a very significant commitment at stake, and the buyer-agent relationship must continue for several weeks or months until a transaction is closed [27]. The relationship dimension usually credited with this is trust in the salesperson. One conception of trust is that it is a composite of the perceived benevolence and credibility of an agent [11]. Other conceptions view trust as a process of uncertainty reduction [1]. In either of these views, trust is something that must be incrementally built up over time, and thus reflects a particular dimension of interpersonal relationship between a salesperson and his or her client.

Within K-6 education, there is evidence that relationships between students are important in peer learning situations, including peer tutoring and peer collaborative learning methodologies [10]. Collaborations between friends involved in these exercises has been shown to provide a

more effective learning experience than collaboration between acquaintances [19]. Friends have been shown to engage in more extensive discourse with one another during problem solving, offer suggestions more readily, are more supportive and more critical than non-friends. In at least one experiment, friends worked longer on the task and remembered more about it afterwards than non-friends.

In psychotherapy, the positive effect of a good therapist-patient relationship on therapeutic outcomes has been demonstrated in several studies, and has even been hypothesized to be *the* common factor underlying the many diverse approaches to psychotherapy which seem to provide approximately equal results [14]. The dimension of the therapist-patient relationship that is credited with these positive attributes is the *working alliance*, which is the non-transfereential bond between the therapist and patient, based on trust and belief in each other as team-members working to achieve the desired outcome.

THE ROLE OF SOCIAL DIALOGUE IN BUILDING RELATIONSHIPS

Given the importance of personal relationships in achieving both social and instrumental objectives, how do people go about building and maintaining them? There is a great deal of evidence that indicates that relationships are primarily developed through social dialogue in the context of face-to-face conversation. Some researchers go as far as to say that relationships themselves are strictly dialogical entities, constituted in talk [12]. Here, I define social dialogue as talk in which the relational aspects of interaction take precedence over instrumental aspects (with "phatic communion" being the purest form of social dialogue [21]). Conversation need not be strictly social *or* task-oriented. The two kinds of conversation can be tightly interleaved in task encounters (for example), and some believe that all messages can be characterized by degree of "phaticity" [9].

Face-to-face interaction is important for relationship building, not only because it is the most primal form of language, but because of the wide range of nonverbal modalities available to interlocutors. Not only are nonverbal skills crucial in the formation of initial impressions in relationships, but they have been shown to correlate with an individual's popularity and the size of their social network [29]. "Immediacy" behaviors (leaning towards one's partner, looking them in the eye, smiling, etc.) have been shown in several studies to be an important channel for communicating affiliation, affect and inclusion with one's partner [22].

RELATIONAL AGENTS

If our software agents are to perform as effective co-workers, sales agents, teachers, and therapists, they need to leverage relationships in the same way that people do, with the ability to intentionally change their relationship with the user as necessary to achieve task goals.

My research in this area has been in the development of *Relational Agents*, which are software agents capable of

building social relationships with people [2]. Relational Agents maintain a representation of the assessed state of the relationship, and use this representation to plan both relationship-appropriate behaviors, as well as behaviors designed to move the relationship along a desired trajectory in order to achieve task objectives.

Relational Agents must maintain a model of their perceived relationship with the user, as a basis for planning relational and task behaviors. Relationships have been modeled in many ways: in terms of the resources they can provide to each partner; in terms of stages of change; in terms of social exchange models (actual equity or equality, investment and commitment); or in terms of dimensions such as power and social distance [3, 4]. I have used the latter characterization, expanding social distance into solidarity (like-mindedness), affect, and familiarity, following Svennevig [30]. I further expand familiarity into a breadth dimension, to capture notions of common ground, and a depth dimension, to capture notions of social penetration theory. This model provides both a fine enough granularity that it can be updated continuously during interaction, and sufficient independent dimensions so that it can represent a wide range of possible relational states.

Relational Agents place many theoretically interesting demands on action planning systems, a number of which have not been adequately – or at all – addressed by existent approaches to planning in classic AI systems. A planner for relational interaction must be able to manage and pursue multiple interactional goals, some or all of which may be persistent or non-discrete. It is not sufficient that the planner work on one goal at a time, since a properly selected utterance can, for example, satisfy a task goal by providing information to the user while also advancing the relational goals of the agent. In addition, many goals, such as intimacy or face goals [4, 9], are better represented by a model in which degrees of satisfaction can be planned for, rather than the discrete all-or-nothing goals typically addressed in AI planners [18]. The discourse planner must also be very reactive, since the user's responses cannot be anticipated. The agent's goals and plans may be spontaneously achieved by the user (e.g., through volunteered information) or invalidated (e.g., by the user changing his/her mind) and the planner must be able to immediately accommodate these changes.

Relational Agents must also be capable of negotiating and participating in a wide range of interactional "frames" [31] including, at a minimum, social and task-oriented frames of interaction. Finally, if situated in the context of face-to-face interaction, Relational Agents must be capable of using and understanding language for both task and relational purposes, as well as recognizing and producing some range of relevant nonverbal conversational and relational cues.

Previous and On-Going Work in Relational Agents

In her book "Life on the Screen", Sherry Turkle describes people's reactions to ELIZA and computer-based

psychotherapy, and concludes that over the last 30 years people have become more comfortable with the idea of computer psychotherapy and relationships with computers, even citing a primitive example of psychoanalytic transference to a computer [33]. In more recent work she has studied people's acceptance of "Relational Objects" such as Tamagotchis, Furbies and robotic dolls which interact with people on a relational and psychological level, "pushing our evolutionary buttons" by synthesizing emotional displays and social behavior.

In a series of studies in the "Computers As Social Actors" paradigm, researchers have demonstrated the possibility of manipulating the user's relationship with a computer using a wide range of behaviors. Reeves & Nass demonstrated that users like computers more when the computer flatters them [28]. Morkes, Kernal and Nass demonstrated that computer agents that use humor are rated as more likable, competent and cooperative than those that do not [25]. Moon demonstrated that a computer that uses a strategy of reciprocal, deepening self-disclosure in its (text-based) conversation with the user will cause the user to rate it as more attractive, divulge more intimate information, and become more likely to buy a product from the computer [24].

Finally, there is a growing body of work on the topic of user trust in computers. Trust is a particularly relevant construct for Relational Agents, since it is a measurable outcome of relationship development. Trust indicates a positive belief about the perceived reliability of, dependability of, and confidence in a person, object, or process, and is one of the key components used in the assessment of computer credibility [32]. Relationally, trust is an antecedent to self-disclosure [34], is an important component of intimacy [1], and trustworthy sources are generally more persuasive [26]. Research on human-computer interfaces has found several interesting results with respect to trust and credibility. It has been found that trust in intelligent systems is higher for systems that can explain and justify their decisions [23]. Also, pedagogical agents, especially those that are highly expressive, affect students' perceptions of trust; these agents are perceived as helpful, believable, and concerned [20].

Small Talk in REA

Embodied Conversational Agents are anthropomorphic computer characters that emulate elements of face-to-face conversation through the use of speech, gaze, gesture, intonation and other nonverbal modalities [8]. Given their use of natural language and nonverbal communicative modalities, ECAs provide an ideal platform for building Relational Agents.

Over the last two years, my work has focused on adding relational capabilities to REA, a real-time, multi-modal, life-sized ECA who has a fully articulated graphical body, can sense the user passively through cameras and audio input, and is capable of speech with intonation, facial

display, and hand gesture [7]. Real estate sales was selected as the application domain for REA because of the opportunity it presents to explore a task domain in which a significant amount of social dialogue normally occurs. Within this domain I have focused on modeling the initial interview between an agent and a prospective buyer, studying the role of social dialog in these encounters, and modeling the behavior of the real estate agent in a dialogue planner for REA [2, 5, 6]. This planner incorporates both task talk—during which REA asks questions about users' buying preferences—and small talk—when REA talks about the weather, events and objects in her shared physical context with the user (e.g., the lab setting), or tells stories about the lab, herself, or real estate.

The planner makes contributions to the conversation in order to minimize the face threat to the user, and maximize trust, while pursuing task goals in the most efficient manner possible. That is, it attempts to determine the face threat of the next conversational move, assess the solidarity and familiarity that currently holds with the user, and judges which topics will seem most relevant and least intrusive to users. As a function of these factors, it chooses whether or not to engage in small talk, and what kind of small talk to choose.

Within this framework, REA decides to do small talk whenever closeness with the user needs to be increased (e.g., before a task query can be asked), or the topic needs to be moved little-by-little to a desired topic and small talk contributions exist that can facilitate this.

In an experiment involving 31 human subjects, small talk was demonstrated to increase users' trust in REA for extroverts (for introverts it had no effect) [6]. One of things I learned from this experiment is that many people feel strongly about their like or dislike of small talk, indicating that either techniques should be developed to automatically determine this preference, or users should be given the ability to move directly into task talk if they desire. I also discovered that purely agent-led social dialogue, without any uptake on user contributions, quickly became unnatural and destroyed the illusion of having a casual chat.

CONCLUSION

We must endow our software agents with relational competencies if they are to meet users' natural expectations, and leverage the power of personal relationships in both social and task-oriented applications. The most promising route to implementing such Relational Agents is through emulating face-to-face conversation with users.

Constructing such agents is an exciting and challenging research area which presents many theoretically interesting problems for future work. For example, exactly how is relational meaning conveyed in social dialogue? It is typically not conveyed in a literal, truth-conditional manner ("I want to be your friend.", "Let's build some rapport."), but indirectly through some form of conversational

implicature [15]. Which elements of relational dialogue are culturally-dependent and which are universal (e.g., as Brown & Levinson claim about Politeness Theory [4])? Which conversational and nonverbal behaviors are especially crucial for different stages of relationship formation and maintenance? How can people guard against deception on the part of relational agents?

My current research is in the area of health behavior change and concerns the development of a personal exercise trainer agent to help users adopt and adhere to a program of regular exercise. In this endeavor I hope to further demonstrate that social dialogue is not just a frivolity to increase the novelty of anthropomorphic agents, but an essential component in leveraging the social psychology of human-computer interaction to achieve outcomes of significant benefit to users.

ACKNOWLEDGMENTS

Thanks to Justine Cassell for advice and guidance in this work, and to Hannes Vilhjálmsón, Dona Tversky and Jennifer Smith for their helpful comments.

REFERENCES

1. Berscheid, E. and Reis, H. Attraction and Close Relationships. *The Handbook of Social Psychology*, D. Gilbert, S. Fiske, and G. Lindzey, (eds.). McGraw-Hill, New York, 1998, 193-281.
2. Bickmore, T. and Cassell, J. Relational Agents: A Model and Implementation of Building User Trust, in *Proceedings of CHI'01*, (Seattle, WA, 2001), 396-403.
3. Brehm, S. *Intimate Relationships*. McGraw-Hill, New York, 1992.
4. Brown, P. and Levinson, S. C. *Politeness: Some universals in language usage*. Cambridge University Press, Cambridge, 1987.
5. Cassell, J. and Bickmore, T. External Manifestations of Trustworthiness in the Interface. *Communications of the ACM* 43, 12, 2000, 50-56.
6. Cassell, J. and Bickmore, T. Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents. *User Modeling and Adaptive Interfaces, in press*.
7. Cassell, J., Bickmore, T., Billinghurst, M., Campbell, L., Chang, K., Vilhjálmsón, H., and Yan, H. Embodiment in Conversational Interfaces: Rea, in *Proceedings of CHI '99*, (Pittsburgh, PA, 1999), 520-527.
8. Cassell, J., Sullivan, J., Prevost, S., and Churchill, E. *Embodied Conversational Agents*. The MIT Press, Cambridge, MA, 2000.
9. Coupland, J., Coupland, N., and Robinson, J. D. "How are you?": Negotiating phatic communion. *Language in Society* 21, 1992, 207-230.
10. Damon, W. and Phelps, E. Strategic Uses of Peer Learning in Children's Education. *Peer Relationships in Child Development*, T. Berndt and G. Ladd, (eds.). Wiley, New York, 1989, 135-157.

11. Doney, P. and Cannon, J. An Examination of the Nature of Trust in Buyer-Seller Relationships. *Journal of Marketing* 61, 1997, 35-51.
12. Duck, S. Talking Relationships into Being. *Journal of Social and Personal Relationships* 12, 4, 1995, 535-540.
13. Gabarro, J. The Development of Working Relationships. *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work*, J. Galegher, R. Kraut, and C. Egido, (eds.). Lawrence Erlbaum Associates, Hillsdale, New Jersey, 1990, 79-110.
14. Gelso, C. and Hayes, J. *The Psychotherapy Relationship: Theory, Research and Practice*. John Wiley and Sons, New York, 1998.
15. Grice, P. *Studies in the Way of Words*. Harvard University Press, Cambridge, MA, 1989.
16. Gumperz, J. Sociocultural Knowledge in Conversational Inference. *Linguistics and Anthropology*, M. Saville-Troike, (ed.). Georgetown University Press, Washington DC, 1977, 191-211.
17. Gutek, B., Cherry, B., Bhappu, A., Schneider, S., and Woolf, L. Features of Service Relationships and Encounters. *Work and Occupations* 27, 3, 2000, 319-352.
18. Hanks, S. Discourse Planning: Technical Challenges for the Planning Community., in *Proceedings of AAAI Workshop on Planning for Inter-Agent Communication*, (1994),
19. Hartup, W. Cooperation, close relationships, and cognitive development. *The company they keep: Friendship in childhood and adolescence*, W. Bukowski, A. Newcomb, and W. Hartup, (eds.). Cambridge University Press, Cambridge, 1996, 213-237.
20. Lester, J. C., Converse, S. A., Kahler, S. E., Barlow, S. T., Stone, B. A., and Bhogal, R. S. The Persona Effect: Affective Impact of Animated Pedagogical Agents, in *Proceedings of CHI '97*, 1997, 359-366.
21. Malinowski, B. The problem of meaning in primitive languages. *The Meaning of Meaning*, C. K. Ogden and I. A. Richards, (eds.). Routledge & Kegan Paul, 1923.
22. Mehrabian, A. Some referents and measures of nonverbal behavior. *Behavioral Research Methods and Instruments* 1, 1969, 213-217.
23. Miller, C. and Larson, R. An explanatory and "argumentative" interface for a model-based diagnostic system, in *Proceedings of the ACM Symposium on User Interface Software and Technology*, (Monterey, CA, 1992),
24. Moon, Y. *Intimate self-disclosure exchanges: Using computers to build reciprocal relationships with consumers*. Harvard Business School, Cambridge, MA Working paper 99-059, 1998.
25. Morkes, J., Kernal, H., and Nass, C. Humor in Task-Oriented Computer-Mediated Communication and Human-Computer Interaction, in *Proceedings of CHI '98*, 1998, 215-216.
26. Petty, R. and Wegener, D. Attitude Change: Multiple Roles for Persuasion Variables. *The Handbook of Social Psychology*, D. Gilbert, S. Fiske, and G. Lindzey, (eds.). McGraw-Hill, New York, 1998, 323-390.
27. Prus, R. *Making Sales: Influence as Interpersonal Accomplishment*. Sage, Newbury Park, CA, 1989.
28. Reeves, B. and Nass, C. *The Media Equation*. Cambridge University Press, Cambridge, 1996.
29. Riggio, R. Social Interaction Skills and Nonverbal Behavior. *Application of Nonverbal Behavioral Theories and Research*, R. Feldman, (ed.). Lawrence Erlbaum, Hillsdale, NJ, 1992, 3-30.
30. Svennevig, J. *Getting Acquainted in Conversation*. John Benjamins, Philadelphia, 1999.
31. Tannen, D. What's in a Frame? Surface Evidence for Underlying Expectations. *Framing in Discourse*, D. Tannen, (ed.). Oxford University Press, New York, 1993, 14-56.
32. Tseng, S. and Fogg, B. J. Credibility and Computing Technology. *Communications of the ACM* 42, 5, 1999, 39-44.
33. Turkle, S. *Life on the screen : identity in the age of the Internet*. Simon & Schuster, New York, 1995.
34. Wheelless, L. and Grotz, J. The Measurement of Trust and Its Relationship to Self-Disclosure. *Human Communication Research* 3, 1977, 250-257.