

Empirical Validation of an Accommodation Theory-based Model of User-Agent Relationship

Timothy Bickmore, Daniel Schulman

College of Computer and Information Science, Northeastern University
360 Huntington Ave - WVH 202, Boston MA 02115
{bickmore,schulman}@ccs.neu.edu

Abstract. We describe a computational model of user-agent relationship based on accommodation theory, in which classes of relationship are defined by the set of activities the user is willing to perform with an agent. An implementation of this model is described that uses dialogue acts as the set of relationship-defining activities, and manipulations of the model to increase user-agent intimacy over time. The implementation is integrated into a virtual agent that plays the role of an exercise counselor. Results from validation studies indicate that the implementation is successful at adapting to users' desired intimacy level, but is not successful at increasing intimacy within the duration of the studies.

Keywords: Relational agent, embodied conversational agent, accommodation theory, personal relationships.

1 Introduction

A growing body of research has investigated the use of appropriate social behavior for virtual agents in different tasks, contexts of use and types of user-computer relationship. However, few of these efforts have made the representation of the user-computer relationship explicit, and even fewer have investigated models in which such relationships can change over time. This is appropriate, since most social agents are only designed to be used in one or a few interactions with a user, and are developed for tasks in which there is little need or norm for changing their relationships over time. However, developing explicit models of user-computer relationship and processes of change is important for applications in which users will be interacting with an agent over very long periods of time or in which explicit management of the relationship is important for task outcomes. Examples of such task domains include sales, education, psychotherapy, chronic disease management, and health behavior change [1].

A significant amount of research has been done in the field of the social psychology of personal relationships on modeling human-human relationships, and these can be used as the basis for representations of human-computer relationships for virtual agents [2]. Examples that have already been integrated into virtual agents include dimensional models comprised of one or more orthogonal scalar measures

such as power and social distance [2, 3], and stage models comprised of a set of well-defined relationship types with an social distance-based ordering among them [4]. Additional models in the social psychology literature include: provision models [5], in which relationships are defined by what they provide to each partner; and economic models [6] consisting of equations describing the relationship between costs, benefits, investment in, and commitments to a relationship.

While these models provide a good starting place, they do not represent the level of detail required for more fine-grained relational planning and reasoning. Ultimately, some representation of the individual beliefs and intentions of the relational partners that underlie these more general models must be used. As social agents become more sophisticated, live longer with their users, and become more sensitive to the social dynamics of their interactions with users, such models will be required to capture the subtleties involved.

In this paper we describe a model of user-agent relationship that represents an advance in this direction (originally proposed in [1]), along with an initial implementation and validation in three longitudinal empirical studies.

2 Related Work

A number of researchers have developed virtual agents that explicitly represent and manipulate their personal relationships with users.

REA. REA is a life-sized Embodied Conversational Agent that played the role of a real estate agent, and her dialog planner modeled the initial interview between a buyer and an agent [3]. The planner dynamically decided between social dialog moves (small talk) and task moves (asking questions about the user’s housing needs) based on an assessment of the current relationship with the user, the face threat of the next desired task move, and several other factors. The relationship was represented using a dimensional model (from [7]), in which solidarity and familiarity are represented as scalars and updated based on the number and content of conversational moves.

FitTrack. The goal of the FitTrack system was to investigate long-term user-computer relationship maintenance in the context of a health behavior change application [2]. The agent in this system plays the role of an exercise advisor that users interact with daily for a month. The agent uses a wide range of techniques from the social psychology of personal relationships—including meta-relational communication, empathy, social dialog, increasing common ground, and nonverbal immediacy behaviors—to establish an increasingly close social bond with the user over the month-long intervention. While these behaviors were intentionally manipulated for the purpose of the intervention study, they were not dynamically planned by the agent. Instead, they were encoded into the agent’s finite-state-machine-based dialogs according to a pre-defined schedule (e.g., the number of conversational turns of social dialog per day). Thus, the relational model was implicitly represented by time (number of interactions with the user).

Autom. The Autom robot is designed to play the role of a weight loss counselor, placed in users’ homes for a longitudinal intervention. The robot represented its relationship with a user with a three-component stage-based model (“acquaintance”,

“relationship buildup”, “relationship maintenance”), and use a limited set of dialogue acts for establishing and repairing its relationship [4].

3 An Accommodation Theory-based Model of Personal Relationships

The relational models used in the systems above were very crude generalizations of the specific beliefs and intentions that an agent and a user have about each other at each point in time. In order to support more nuanced reasoning and planning by an agent, a more sophisticated approach is required.

Models of multi-agent collaboration (such as SharedPlans [8]) provide a potential starting place. Examples of such relational collaborations involve coordination on specific activities within a relationship (e.g., washing and drying the dishes, reminiscing) as well as collaboration on the relationship itself (e.g., negotiating roles). However, theories of multi-agent collaboration generally are concerned with the accomplishment of a specific goal, using specific actions over a specific time interval. Relationships, on the other hand, are typically unbounded in duration, and while the range of activities conducted within the relationship can be specified, the particular activities that a dyad engages in at any one time cannot be defined. Further, while specific actions may be required to build, change, maintain or terminate a relationship, no actions are required to simply “have” a relationship (e.g., partners can say they’re friends even if they haven’t talked to each other in ten years).

Thus, rather than a collaboration, per se, we model a relationship as the set of tasks that two agents (or an agent and a user) are ready and willing to collaborate on at any given time. This notion can be defined in terms of Thomason’s Accommodation Theory, in which accommodation is defined as the situation in which one agent infers the goals of a second and takes action to help without the first agent making an explicit request [9]. Defining relationships based on the set of activities that two people routinely perform together is also the basis of provision-based relational models in social psychology (e.g., [5]).

People rarely talk explicitly about relationship (e.g., “Do you want to be my friend?”), but rather infer relational status from actions and what their partner is willing to do with them. Thus, we are talking about modeling actions at two levels: (L1) the specific observable actions that two agents explicitly negotiate and perform; and (L2) the relationship status. The relationship status is typically not negotiated explicitly, but rather through the actions at L1, and thus requires accommodation to make relational inferences. In addition, given a known relationship status, actions at L1 can be accommodated without negotiation.

This general model can be used: 1) to assess the current status of the user’s perception of the their relationship with an agent by observing what relational actions they agree to or initiate; 2) to move the relationship in a desired direction, by proposing actions from relationship categories different from the current one and assessing the user’s response; and 3) to govern the behavior of the agent in collaborating on actions in the current relationship category.

3.1 Implementation of the Model

The accommodation-based model of relationship was implemented for a virtual agent that provides longitudinal health counseling to users [10]. Although the theory is very general regarding the kinds of actions that agents can perform, in our initial implementation all agent actions are dialogue acts that the user can choose to participate in or not. We also use a single ordinal variable to represent social distance, which we refer to as “intimacy” (more general than the romantic kind). Each intimacy value indexes a set of dialogue actions that are appropriate for that kind of relationship. Our implementation leverages the theory by making the assumption that accommodation on one act within a relationship category (specified by the intimacy value) implies accommodation on others.

The implementation is based on a specific set of relational dialogue acts that were pre-sorted into four categories of relationship, by increasing intimacy, appropriate for a health counselor:

0. *stranger/professional* - things you would expect a fairly impersonal counselor to say on your first encounter.
1. *more than a professional relationship* - can go into casual off-topic chat, somewhat personal, also meta-relational talk about the working relationship.
2. *casual friends* - talk about anything but the most intimate topics.
3. *close friends* - almost anything goes, including love life, near-death experiences, embarrassing moments, etc.

A set of 109 agent relational dialogue acts were authored, based on an experimental manipulation for increasing interpersonal closeness between people in the laboratory [11], and from review of transcripts of actual exercise trainer/client dialogues. The dialogue acts were then independently sorted into the four categories by four judges. Intraclass correlation among the judges was significant ($p < .001$), demonstrating reliability in their assessments. A final classification of messages was performed by reconciling disagreements using an averaged score decision rule. Table 1 shows a sample of the relational messages from each category.

We designed a relationship management algorithm for the agent that could operate for multiple interactions with a user over months or years of operation. The two primary design goals were to automatically adapt the agent’s model of relationship to a user, and to incrementally move relationships with users towards increasing intimacy. Increasing intimacy may be important for a health counseling agent, under the assumption that a better working relationship will lead to greater retention in the intervention and higher compliance with the agent’s requests (the quality of client-therapist working relationship has been demonstrated to have a significant correlation with outcome measures across a wide range of problems, therapeutic approaches, and outcome measures [12]).

Table 1. Example Relational Dialogue Acts

Category	Example Dialogue Acts by Agent Counselor
0. Stranger / Professional	“I really appreciate your determination in meeting your health-related goals.”; “Remember my job is to help you, so let me know if you need to talk at any time.”
1. More than Professional	“You know, I hope I can be a stable source of support for you while we work together.”; “I think we both like to talk about your exercise.”
2. Casual Friend	“Can I tell you a secret? You know sometimes I feel a little guilty when I just keep telling you what to do. I’m sorry if I sound too much like a nag sometimes. It’s just that your health is very important to me.”; “Given the choice of anyone in the world, who would you want as a dinner guest?”
3. Close Friend	“Can I confide in you? Sometimes I get scared when we end our conversations. I mean, I just disappear, and maybe I won’t ever come back. It’s a little like death, I think. But then I see you again and I feel great.”; “You know, sometimes I wish that I could cry, but my programmers did not give me that ability. When was the last time you cried in front of someone?”

Adaptation of the relational model is performed using both implicit and explicit assessments of the user’s desired level of intimacy. Implicit assessments—following accommodation theory—are based on user reactions to relational messages offered by the agent. Conversational “uptakes”, or accommodation, provides evidence that the user is comfortable with the level of intimacy implied by the dialogue act. Conversational “rejects”—in which the user explicitly or implicitly indicates they do not want to collaborate on a proposed dialogue act—provides evidence that the user would prefer less intimacy than that implied by the dialogue act. Explicit assessments of relationship are performed using a self-report questionnaire asking users to classify their relationship with the agent into one of the four categories.

As an example, the agent may make a relational bid by saying “Thanks for sharing all of your thoughts and opinions with me. I find them very interesting, and I hope you don’t find it too tedious.”, in which case the user is given the choice of responding with one of the following options: 1) “No, not all.” (uptake); 2) “Sometimes.” (uptake); 3) “Well, since you mention it, I do.” (uptake); or 4) “That seems like an inappropriate question.” (rejection).

The agent attempts to increase intimacy with a user over time by periodically using a dialogue act that is slightly more intimate than the intimacy level of the current relationship. As described above, this intimacy “bid” can be rejected by a user (causing intimacy to remain at its current level), or they can uptake on the offer (causing modeled intimacy to increase to the next level).

Relational Algorithm. The relational algorithm maintains a single **intimacy** variable, ranging over the four values described above, initialized at 1 (“More than Professional”) for a new user. During each counseling conversation with a user (up to once daily in the health counseling application) the agent makes a “bid” to engage the user in one relational dialogue act by uttering the relational message at an appropriate point in the conversation, then allowing the user to “uptake” (accommodate) the act or reject it, through dialogue actions. Note that at intimacy level 0 the agent still makes

relational bids, however, these represent dialogue acts that strangers would typically use with each other.

The following rules describe the longitudinal algorithm:

- Every D days (or nearest interaction afterwards) there is an explicit assessment of intimacy, instead of a relational message. This is performed by displaying a text-based questionnaire for the user to respond to at the end of the session after the agent has performed a farewell. The **intimacy** variable is always set to the user's assessment.
- If there have been B relational messages delivered since the last change in **intimacy** (for any reason), the agent makes a "bid" to increase intimacy, by selecting a relational dialogue act from the **intimacy+1** category (if any). Otherwise, a relational dialogue act is selected from the **intimacy** category and initiated.
 - If the user's response is a rejection, and the dialogue act is not a bid, then **intimacy** is decremented.
 - If the user's response is an uptake, then **intimacy** is incremented.

The parameters D and B varied across the three studies described here: $D=30$ and $B=5$ in studies 1 and 2, while $D=7$ and $B=3$ in study 3, below.

Virtual Agents. In the systems used for experimentation with this algorithm, users have (up to) a daily 10 minute conversation with the virtual exercise counselor on their home computer. Dialogues are scripted using a custom hierarchical transition network-based scripting language. Agent nonverbal conversational behavior is generated using BEAT [13], and includes beat (baton) hand gestures and eyebrow raises for emphasis, gaze away behavior for signaling turn-taking, and posture shifts to mark topic boundaries, synchronized with synthesized speech. User input is obtained via multiple choice selection of utterances [2]. The dialog used in the non-relational portion of the counseling conversation is similar to that described in [14, 15].

4 Preliminary Validation Studies

We have two primary research questions in all of the following empirical studies. First, is user relational behavior (specifically uptakes and rejections of agent relational bids) a good way to assess user-agent relationship? We want to determine whether the accommodation model provides a reliable and valid assessment mechanism for user-agent relationship, as evidenced by longitudinal stability in the intimacy assessments (test-retest reliability) and correlation with other measures of relationship (convergent construct validity). Second, is the relational algorithm described above effective at building relationships? We want to determine whether making incremental bids to increase intimacy improves the user-agent relationship over time, as evidenced by increasing intimacy and comparison to agents that are not using this model.

Common Measures. We logged **intimacy** variable values, whether users logged in on a given day or not, and the periodic explicit user assessment of intimacy (as

described above). We also periodically asked users to assess their working alliance bond with the agent and a single scale item reflecting desire to continue working with the agent, both administered via text forms at the end of a session, after the agent has concluded its conversation with the user. Working Alliance reflects the trust and belief that a client has in working with a helper to achieve a desired therapeutic outcome. The Working Alliance construct has three sub-components: a goal component, reflecting the degree to which the helper and client agree on the goals of the therapy; a task component, reflecting the degree to which the helper and client agree on the therapeutic tasks to be performed; and a bond component, reflecting the trusting, empathetic relationship between the client and helper. We assessed user-reported Working Alliance using items from the revised short form of the Working Alliance Inventory [16]. The number of steps walked each day was also recorded, as measured by an Omron pedometer and electronically uploaded to the system to support the exercise intervention.

Study 1. Our initial validation of the model was done in the context of an exercise promotion agent used by sedentary older adults [10]. In this implementation, when users were presented with relational messages, they were given the choice of uptake by selecting an utterance that indicated a desire to engage in the dialogue act, or reject by selecting an unambiguous “bald” statement of rejection (e.g., “I’m not comfortable talking about that with you.”).

A total of 32 participants, age 55 or older, took part in the study. Participants were randomized to receive the exercise promotion counseling intervention with or without the inclusion of the relational model described above. Participants were in the study for 159 days on average (range 9-191), conducting a total 5,160 interactions with the agent. At the end of this time there were no significant differences between groups on measures of relationship closeness. The most striking finding, however, was that in 291 relational bids there was not a single instance of user rejection.

Study 2. An identical relational model was integrated into a conversational agent that provided a year-long health behavior change intervention designed to promote exercise and UV (sunlight) avoidance [17]. At the time of this writing, 181 participants had conducted 997 interactions with the agent, without a single instance of a user rejection of a relational bid.

Discussion. There were several possible reasons why the relational model failed to lead to increases in intimacy in Study 1, as evidenced by users not rejecting a relational bid, even after several hundred trials. First, users may not be treating their interaction (and relationship) with the agent as seriously as they do their interactions with other people, so assumptions about modeling user-agent relational behavior after human-human ones may not be valid (i.e., they may just be selecting responses at random, or to see “what happens”). However, another possible reason why the relational model did not perform as expected is that the model did not follow human behavior closely enough, at least in terms of rejections. In human-human interactions, bids for increased intimacy are almost never baldly rejected, but declined in ways that save the requestor’s positive face [18]. For example, if a new acquaintance were to suddenly start asking you about your love life, you would likely find an excuse to end

the conversation or change the topic, rather than telling them that their inquiries are inappropriate.

5 Final Validation Study

Under the assumption that the primary problem with the first two studies was with lack of fidelity, we extended the relational model so that users were given a range of responses to each relational bid, reflecting varying degrees of uptake or rejection. For example, for the relational act “So, do you have a lot of friends?”, rather than letting users only select “Yes”, “No”, or “That’s inappropriate.”, we now let them select from responses designed to span the range from eager uptake to hedged desire to change the topic, in addition to removing the extreme rejection option (“That’s inappropriate.”) since we established that users would never select it:

Agent: “So, do you have a lot of friends?”
User: “Let's talk about my friends.” (uptake), *or*
“Yes, I do.” (affirmative response), *or*
“No, I don't.” (negative response), *or*
“I guess.” (ambiguous or hedged response) *or*
“Let's talk about my walking.” (rejection)

In order to determine whether these relational bid responses can be reliably interpreted as being ordered from a strong uptake of the bid to a strong rejection, we conducted a pilot study to test ratings. Ten participants were shown six relational bids and asked to rate the bid responses along a strong reject to strong uptake scale. Results indicated unambiguous meanings for the uptake and rejection moves above, with the other moves in-between the extremes but largely unordered.

Experimental Methods. This study used the same virtual exercise promotion framework used in Study 1 [10]. Aside from allowing users to express a wider range of responses to relational dialogue acts, the model and study protocol were identical to those in Study 1.

Participants. Twenty-one subjects, aged 55-70, 95% female participated in the study. All had been using the system, and interacting with the agent, for a minimum of two months prior to this study. Eleven were randomized into the relational condition and 10 into a control group who had identical interactions with the agent except that the relational model described above was not used (no relational messages).

Aggregate Results. All statistics were calculated using R-2.15.0¹. The study was active for 37 days. Participants used the system on 499 of 653 opportunities (76.4%).

Relational Responses. Analysis of user responses to relational messages indicates that the full range of dialogue responses were used, but that use of uptake and reject responses decreased over time (Figure 1).

¹ <http://www.R-project.org>.

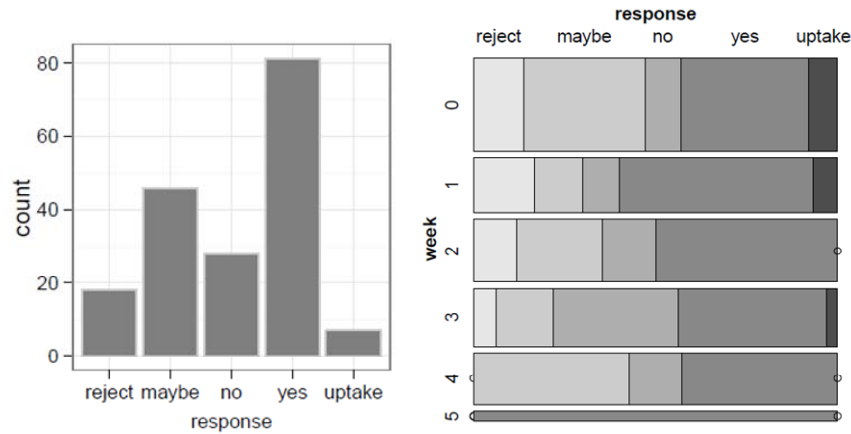


Fig. 1. User Responses to Relational Messages
(Left: total frequency; Right: proportions over time)

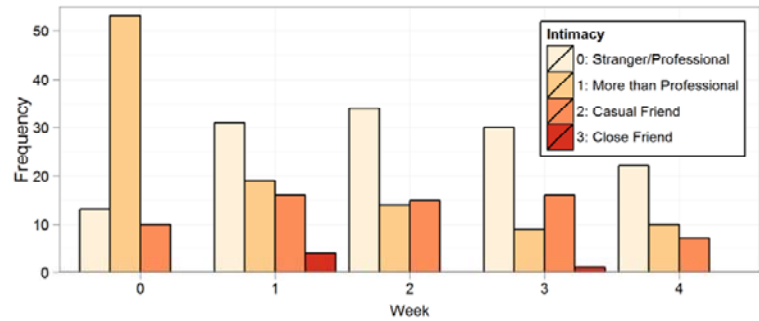


Fig. 2. Changes in Intimacy Variable Over Time

Intimacy Change Over Time. The most common value for the intimacy variable is '1' in week 1 (its initial value). After this, however, the most frequent value decreases to '0' where it remains (Figure 2). This is also reflected by fitting a mixed-effect ordinal logistic regression of intimacy on linear and quadratic effects of study day: modeled intimacy decreases significantly from the start of the experiment (linear effect; $b=-0.103$, $SE=0.042$, $p=0.016$), but this decrease levels off over time (quadratic effect; $b=0.002$, $SE=0.001$, $p=0.047$), demonstrating longitudinal stability.

Relationship between Intimacy and other Relational Measures. The relationship between intimacy and other measures was tested with a series of mixed-effect ordinal logistic regressions of intimacy on each measure as a predictor. There was a significant relationship between modeled intimacy and the Working Alliance bond measure ($b=0.867$, $SE=0.181$, $p<0.001$). There was a near-significant relationship between intimacy and self-reported desire to continue working with the agent (likelihood ratio $\chi^2(4)=8.51$, $p=0.075$). However, there was no significant relationship between modeled intimacy and actual system use ($b=-0.018$, $SE=0.349$, $p=0.96$).

Relationship between Intimacy, System Use and Outcomes. There was no significant correlation between intimacy level and exercise (mixed effect model), nor

system use and exercise (mixed effect model aggregating by week), although the latter did indicate a trend in the expected direction.

Between-Group Comparisons. There were no significant differences between the relational and control groups on any relational measure, including explicit user intimacy assessments, working alliance bond, self-reported desire to continue working with the agent, and system usage (frequency of logins).

Case Studies. Understanding complex longitudinal models is often best accomplished through detailed analysis of individual cases. Here we look at three users in the relational group.

Participant #5 – Rapid Intimacy. 64 year-old female. (Figure 3) Although #5 was only in the study for 9 days (late starter), her experience exemplifies agent adaptation to a user who is ready and willing to increase intimacy. On days 3 and 6 she responded to relational messages with explicit uptakes, causing her intimacy level to increment up to its maximum level. She completed an explicit assessment of intimacy on her final day, indicating it was at level 2 – generally consistent with the model – but would have caused her intimacy to revert to level 2 had she logged in again. She reported her alliance bond score and desire to continue consistently at their maximum values, both consistent with her behavior.

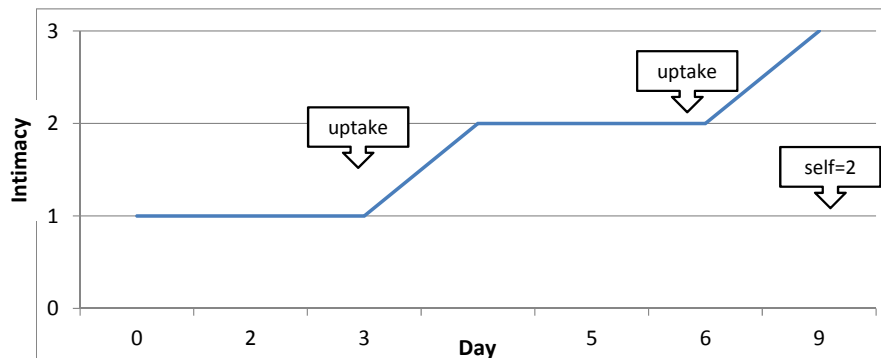


Fig. 3. Case Study for Participant #5

Participant #25 – Zero Intimacy. 62 year-old female. (Figure 4). This participant exemplifies a user who does not want any intimacy with the agent and is unchanging in her attitude. She uses explicit rejects four times in response to relational messages from the agent, and consistently self-reports her intimacy at level 0. She reported her alliance bond score and desire to continue consistently at near-minimum values, all consistent with her behavior.

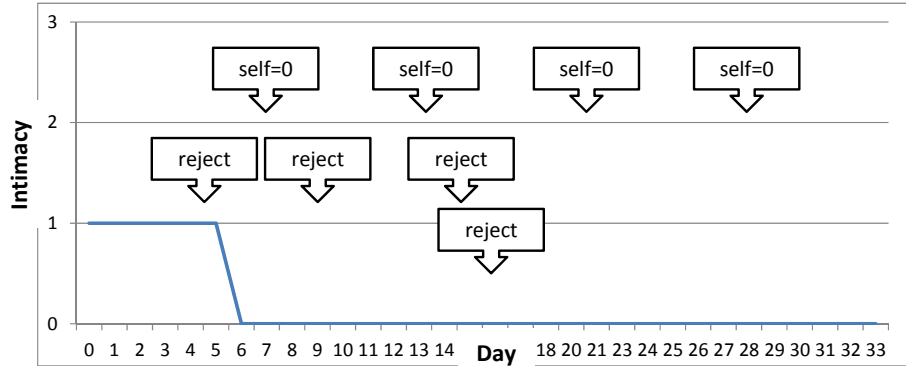


Fig. 4. Case Study for Participant #25

Participant #27 – Inconsistent. 57 year-old female. (Figure 5). This participant provided inconsistent self-reports of intimacy, spanning the full range of values, while also frequently rejecting the agent’s relational messages (8 times in total), causing the intimacy variable to oscillate wildly over time. She reported working alliance and desire to continue at maximum values throughout the month. This was the only one of the 11 intervention participants with this pattern of behavior, although we have observed similar intimacy oscillations in Studies 1 and 2.

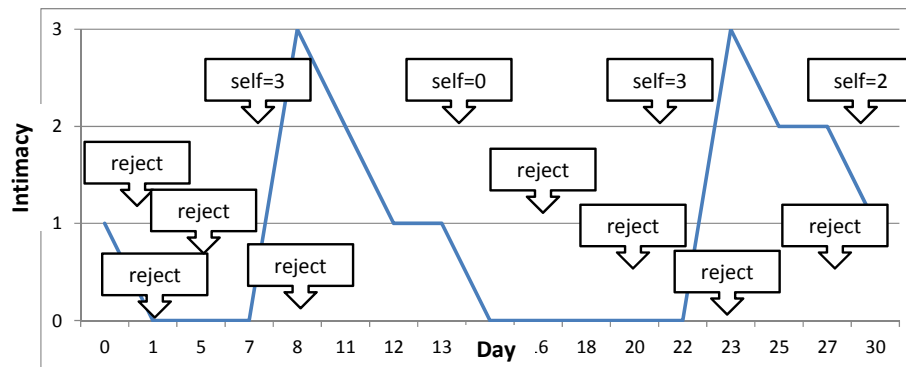


Fig. 5. Case Study for Participant #27

Discussion. Regarding our original research questions, we did find that the accommodation model provides a reliable assessment mechanism for user-agent relationship. We find that intimacy levels off over time, both through statistical tests of the aggregate data and from visual inspection of individual cases (it is stable after 2 weeks for 6 of the 11 participants). We also found that the accommodation model provides a valid assessment mechanism of user-agent relationship: convergent construct validity was established between modeled intimacy and the Working Alliance and desire to continue relational measures. However, we did not find that the accommodation model, at least in its current implementation, leads to increased

intimacy over time, based on comparisons between the relational and control groups or based on trend analysis, which demonstrated that, if anything, intimacy decreases over time.

One limitation of this study is that we did not explicitly control for relational message length. An analysis of the 109 messages (exemplified in Table 1) using a one-way ANOVA indicates that, while the differences are not statistically significant, there is a clear trend towards increasing message lengths in relational categories 2 and 3. Upon further analysis, this is due to increased agent self-disclosure in the middle categories.

5 Conclusions and Future Work

We presented a series of studies attempting to validate an accommodation-based model and implementation of user-agent relationship. The final study provided evidence that the model can be used to reliably and validly assess the status of user-agent relationship and to drive the use of relationship-appropriate activities. We were unable to show that this model can be used to gradually nudge users into more intimate relationships with an agent, at least within the timeframe of the studies.

There are many reasons why the model failed to improve intimacy. As stated in Section 4, two reasons may be that users do not treat their relationship with an agent as they do their relationships with other people (perhaps not even taking the relational messages seriously), or that our model and implementation still lack enough fidelity. Human relational negotiation is very complex, since it is usually conducted in a tacit, off-record manner, with the bids, uptakes, and rejections handled in an indirect manner to prevent explicit rejection and loss of self-esteem. The communication channel between users and our agent (via multiple choice menu input) may be much too impoverished to support this subtlety.

Another reason for the lack of increasing intimacy may simply be that the rules or parameters we used in our model need to be changed (e.g., the frequency of bids by the agent to increase intimacy).

Our model may also simply be inappropriate for some users, who have fundamentally different ideas about the kinds of activities that are appropriate for different types of relationship (e.g., due to cultural differences). One explanation for the behavior of subject #27 above may be that she fundamentally disagreed with our pre-ranking of relational acts that are appropriate for a professional health counselor.

Future Work. There are many interesting directions of future research. Following Argyle and Dean's equilibrium theory of intimacy, unwanted advances in intimacy (bids) in one communication channel (such as speech) can result in compensation behaviors in other channels (such as proxemics or gaze) [19]. Thus, in addition to providing users with more nuanced and hedged forms of bid rejection in their speech, they may also make use of their real or simulated nonverbal behavior, for example moving themselves away from the agent in the real or virtual world.

Intimacy may also be improved through a greatly expanded repertoire of relational dialog, including such things as more in-depth reciprocal self-disclosure exchanges to establish more common ground with the user.

Another future area of exploration is giving users the ability to make relational bids. Although many may not use this function, for some (such as subject #5 above) it would allow them to negotiate their desired intimacy level much more quickly than our current algorithm allows. It may also be more effective when a change in working relationship is the user's idea instead of the agent's.

Our relational model is admittedly over-simplistic, and development of more sophisticated models represents another important direction of research. The number of intimacy levels modeled and the sets of relational actions indexed should both be functions of context and user role. Ultimately, relational models based on specific beliefs and intentions of the agent and user will be needed for maximum fidelity.

A final note on methodology. While longitudinal models of user-agent relationship represent an open and promising area of research, studies of actual user responses to such models take very long periods of time to conduct (the studies described here took over two years to complete). The development of evaluation and simulation tools that can provide insights into longitudinal human behavior in shorter time frames is thus another important area of investigation.

Acknowledgments. Thanks to Barbara Barry and Laura Pfeifer for comments on early drafts, and to the anonymous reviewers for their many suggestions. This material is based upon work supported by the National Science Foundation under Grant No. 0545932. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

References

1. Bickmore, T.: Relational Agents: Effecting Change through Human-Computer Relationships. Ph.D. Dissertation, Media Arts & Sciences. Massachusetts Institute of Technology, Cambridge, MA (2003)
2. Bickmore, T., Picard, R.: Establishing and Maintaining Long-Term Human-Computer Relationships. *ACM Transactions on Computer Human Interaction* 12, 293-327 (2005)
3. Cassell, J., Bickmore, T.: Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents. *User Modeling and Adaptive Interfaces* 13, 89-132 (2003)
4. Kidd, C.D.: Designing Long-Term Human-Robot Interaction and Application to Weight Loss. Ph.D. Dissertation, Media Arts & Sciences. Massachusetts Institute of Technology, Cambridge, MA (2008)
5. McGuire, A.: Helping Behaviors in the Natural Environment: Dimensions and Correlates of Helping. *Personality and Social Psychology Bulletin* 20, 45-56 (1994)

6. Rusbult, C., Drigotas, S., Verette, J.: The investment model: An interdependence analysis of commitment processes and relationship maintenance phenomena. In: Canary, D., Stafford, L. (eds.) *Communication and relational maintenance*, pp. 115-139. Academic Press, San Diego (1994)
7. Svennevig, J.: *Getting Acquainted in Conversation*. John Benjamins, Philadelphia (1999)
8. Grosz, B., Kraus, S.: Collaborative plans for complex group action. *Artificial Intelligence* 86, 269-357 (1996)
9. Thomason, R.: Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics. In: Cohen, P., Morgan, J., Pollack, M. (eds.) *Intentions in Communication*, pp. 325-364. MIT Press, Cambridge, MA (1990)
10. Bickmore, T., Schulman, D.: A Virtual Laboratory for Studying Long-term Relationships between Humans and Virtual Agents. In: *Proceedings Autonomous Agents and Multi-Agent Systems, Budapest (2009)*
11. Aron, A., Melinat, E., Aron, E.N., Vallone, R.D., Bator, R.J.: The experimental generation of interpersonal closeness: A procedure and some preliminary findings. *Personality and Social Psychology Bulletin* 23, 363-377 (1997)
12. Horvath, A., Symonds, D.: Relation Between Working Alliance and Outcome in Psychotherapy, A Meta-Analysis. *Journal of Counseling Psychology* 38, 139-149 (1991)
13. Cassell, J., Vilhjálmsón, H., Bickmore, T.: BEAT: The Behavior Expression Animation Toolkit. In: *Proceedings SIGGRAPH '01*, 477-486 (2001)
14. Bickmore, T., Caruso, L., Clough-Gorr, K., Heeren, T.: "It's just like you talk to a friend" - Relational Agents for Older Adults. *Interacting with Computers* 17, 711-735 (2005)
15. Bickmore, T., Gruber, A., Picard, R.: Establishing the computer-patient working alliance in automated health behavior change interventions. *Patient Education & Counseling* 59, 21-30 (2005)
16. Hatcher, R.L., Gillaspay, J.A.: Development and validation of a revised short version of the working alliance inventory. *Psychotherapy Research* 16, 12-25 (2005)
17. Velicer, W., Bickmore, T., Byron, D., Johnson, J.: Using Relational Agents in Interventions for Multiple Risk Factors. In: *Proceedings of the Society for Behavioral Medicine Annual Conference*, (2011)
18. Goffman, I.: *On face-work*. *Interaction Ritual: Essays on Face-to-Face Behavior*, pp. 5-46. Pantheon, New York (1967)
19. Argyle, M., Dean, J.: Eye-contact, distance and affiliation. *Sociometry* 28, 289 – 304 (1965)