

**Relational Agents:
Effecting Change through Human-Computer Relationships**

by
Timothy Wallace Bickmore

Master of Science, Computer Science,
Arizona State University, 1986
Bachelor of Science and Engineering, Computer Systems Engineering
Arizona State University, 1985

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
In partial fulfillment of the requirements for the degree of

Doctor of Philosophy
at the
Massachusetts Institute of Technology
February 2003

© 2003 Massachusetts Institute of Technology, All rights reserved.

Author

January 31st, 2003

Certified by

Rosalind W. Picard
Associate Professor of Media Arts and Sciences
Thesis Supervisor

Certified by

Justine Cassell
Associate Professor of Media Arts and Sciences
Thesis Supervisor

Accepted by

Andrew B. Lippman
Chairperson, Departmental Committee on Graduate Students

Relational Agents: Effecting Change through Human-Computer Relationships

by
Timothy Wallace Bickmore

Submitted to the Program in Media Arts and Sciences,
School of Architecture and Planning,
In partial fulfillment of the requirements for the degree of
Doctor of Philosophy
at the
Massachusetts Institute of Technology

Abstract

What kinds of social relationships can people have with computers? Are there activities that computers can engage in that actively draw people into relationships with them? What are the potential benefits to the people who participate in these human-computer relationships?

To address these questions this work introduces a theory of Relational Agents, which are computational artifacts designed to build and maintain long-term, social-emotional relationships with their users. These can be purely software humanoid animated agents--as developed in this work--but they can also be non-humanoid or embodied in various physical forms, from robots, to pets, to jewelry, clothing, hand-helds, and other interactive devices. Central to the notion of relationship is that it is a persistent construct, spanning multiple interactions; thus, Relational Agents are explicitly designed to remember past history and manage future expectations in their interactions with users. Finally, relationships are fundamentally social and emotional, and detailed knowledge of human social psychology--with a particular emphasis on the role of affect--must be incorporated into these agents if they are to effectively leverage the mechanisms of human social cognition in order to build relationships in the most natural manner possible.

People build relationships primarily through the use of language, and primarily within the context of face-to-face conversation. Embodied Conversational Agents--anthropomorphic computer characters that emulate the experience of face-to-face conversation--thus provide the substrate for this work, and so the relational activities provided by the theory will primarily be specific types of verbal and nonverbal conversational behaviors used by people to negotiate and maintain relationships.

This work also provides an analysis of the types of applications in which having a human-computer relationship is advantageous to the human participant. In addition to applications in which the relationship is an end in itself (e.g., in entertainment systems), human-computer relationships are important in tasks in which the human is attempting to undergo some change in behavior or cognitive or emotional state. One such application is explored here: a system for assisting the user through a month-long health behavior change program in the area of exercise adoption. This application involves the research, design and implementation of relational agents as well as empirical evaluation of their ability to build relationships and effect change over a series of interactions with users.

Thesis Supervisors: Rosalind W. Picard, Associate Professor of Media Arts and Sciences
Justine Cassell, Associate Professor of Media Arts and Sciences

Doctoral Dissertation Committee

Thesis Co-Advisor

Rosalind W. Picard, Sc.D.
Associate Professor of Media Arts and Sciences
M.I.T. Program in Media Arts and Sciences

Thesis Co-Advisor

Justine Cassell, Ph.D.
Associate Professor of Media Arts and Sciences
M.I.T. Program in Media Arts and Sciences

Thesis Reader

Amanda J. Gruber, M.D.
Assistant Professor of Psychiatry
Harvard Medical School

Thesis Reader

Candace L. Sidner, Ph.D.
Senior Research Scientist
Mitsubishi Electric Research Laboratory

January 2003

Acknowledgements

Thanks, first and foremost, to a wonderfully supportive and helpful committee. Roz provided the initial inspiration that led me into the field of health behavior change, and I could not have pulled this project off without her having supported my work through to its completion. My discussions with Candy--ranging from philosophy to linguistics to artificial intelligence--were always insightful, and it has been a real pleasure to interact with someone who cares deeply about both theoretical issues *and* building useful systems. Amanda gave generously of her time and knowledge, even before I asked her to serve on my committee, and provided me with the crucial knowledge of how various therapeutic approaches to change are applied in a clinical setting. Justine taught me to look at human behavior in a whole new way.

Thanks also to behavior change researchers at the Cancer Prevention Research Center at the University of Rhode Island, especially Wayne Velicer who helped me develop the experimental design for the behavior change study. Bryan Blissmer and Deb Riebe gave me essential tips and feedback on how to structure an exercise adoption study.

My work was supported by the many wonderful sponsors of the MIT Media Lab and NSF ROLE grant REC-0087768.

Development of the FitTrack system and study was a huge undertaking, involving lots of people over the year or so it was in development. Ian Gouldstone was designer and animator extraordinaire, who designed the characters used in the system, and consulted on various aspects of the system's design. Marlisa Febbriello of the Health Education department at MIT Medical provided invaluable help in developing the educational content used in the system. Julie Stalma and Dorri Lee were the exercise trainers who acted as role models for the exercise advisor character. Laura Packer wrote many of the exercise stories used in the system. Nancy Alvarado of IBM Research provided useful feedback on the facial expressions used by the character. Anna Pandolfo did most of the subject intake interviews, and provided feedback on various aspects of the experimental design. Sherry Turkle provided great suggestions for the qualitative interviews done at the end of the study. Dan Ariely provided invaluable feedback on the statistical analysis of data from the study. Noah Fields and Jeannie Finks of the Media Lab's NeCSys group provided many helpful suggestions and software and support for the web server and relational database used in the study. Karen Liu helped design some of the empathy dialogue networks in the system. Jonathan Klein provided lots of useful feedback on the system design. Thanks to Mindy Chang, Chris Cary and Temitope Sibun--undergraduate researchers--who put many pieces of the system in place over the Summer of 2002, and to Tom Stocky for patiently tutoring them. And, thanks to the one hundred-plus study participants who proved that it was all worthwhile.

The REA small talk study also involved a number of other people, including all the past developers of the REA system (notably Hannes Vilhjálmsson, Lee Campbell and Hao Yan), and folks who contributed to the design, implementation and execution of parts of the experiment, including Jennifer Smith, Elisabeth Sylvan, and Ian Gouldstone (as the "wizard").

Thanks to Linda Peterson and Pat Solakoff for essential advice on the mechanics of the Ph.D. program, and for their support through turbulent times.

A warm thanks to my work buddies and moral support network for the last four-and-a-half years, Hannes Vilhjálmsson and Kimiko Ryokai. Thanks also for the support and

feedback from past and current members of the Gesture and Narrative Language and Affective Computing groups at the Media Lab.

Thanks to my family for pretending to understand why I came back to school after a decade in industry, and to my girls, Katie and Olivia, for giving up their father during my time at MIT.

Finally, my deepest gratitude goes to my partner Jennifer, who provided the emotional and social support I needed to get through this. Not only did she keep me sane during this time, but she played a substantial role as sounding board and proof-reader, and implemented and tested a significant portion of the FitTrack dialogue networks and other parts of the system.

Table of Contents

Abstract.....	3
Doctoral Dissertation Committee	5
Acknowledgements	7
Table of Contents	9
Chapter 1. INTRODUCTION	15
1.1 Motivation: The Importance of Personal Relationships	15
1.2 Relational Agents.....	17
1.3 Embodied Conversational Agents	18
1.4 Two Domains of Analysis	18
1.5 Summary	18
Chapter 2. RELATED WORK.....	21
2.1 Introduction.....	21
2.2 Relationships with Non-relational Entities	21
2.2.1 Trust in Man-Made Artifacts	22
2.3 Social Agents	23
2.4 Affective Agents	23
2.5 Anthropomorphic Agents	24
2.5.1 Embodied Conversational Agents	25
2.5.2 Sociable Robots	26
2.5 Relational Agents.....	27
2.7 Other Related Work	28
2.8 Changing Attitudes Towards Relational Agents	28
2.9 Summary	30
Chapter 3. THEORY	31
3.1 The Social Psychology of Personal Relationships.....	31
3.2 Multi-Agent Collaboration	32
3.3 Accommodation.....	35
3.4 A Formal Definition of Relationship.....	36
3.5 Keeping Score.....	37
3.5.1 Trust.....	38
3.6 Motivation to Meet Relational Expectations	40
3.6.1 Rights & Obligations	40
3.6.2 Implications	41
3.7 Relational Dynamics.....	41
3.7.1 Negotiation.....	41
3.7.2 Face.....	41
3.7.3 Other Forms of Change.....	42
3.8 The Affective Dimension.....	43
3.8.1 Affective Relational Provisions	43
3.8.2 The Regulatory Role of Affect	44
3.8.3 Affect in Relational Negotiation.....	44
3.9 Application to the Micro-Structure of Face-to-Face Conversation	44

3.9.1 The Role of Nonverbal Behavior.....	45
3.9.2 Relational Communication	45
3.9.3 Social Deixis	46
3.9.4 Social Dialogue.....	47
3.9.5 A More General Model of Facework in Face-to-Face Conversation	48
3.9.6 Face-to-Face Conversation with Relational Agents	51
3.10 Application to the Macro-Structure of Long-Term Relationships.....	51
3.10.1 Relationship Maintenance.....	52
3.10.2 Expected Trajectories of Relational Change	53
3.10.3 The Role of Personal History.....	53
3.10.4 Persistent Relational Agents	54
3.11 Summary	54
Chapter 4. DISCOURSE PLANNING FOR RELATIONAL DIALOGUE	55
4.1 Introduction.....	55
4.1.1 Problem: Dialogue Move Sequencing to Achieve Task Goals under Relational Constraints	55
4.1.2 Previous Work in Discourse Planning	56
4.2 Rea	57
4.2.1 Rea Architecture	58
4.2.2 Discourse Planner Interface	59
4.3 Relational Model	60
4.4 Small Talk.....	61
4.4.1 How Small Talk Works	61
4.5 Discourse Planner Design.....	63
4.5.1 User Modeling of Human-Computer Relationships	64
4.5.2 Discourse Planning for Mixed Task and Social Dialog.....	65
4.6 Example Interactions	71
4.7 Summary	75
Chapter 5. EVALUATION OF FACE-TO-FACE INTERACTION WITH A RELATIONAL AGENT	77
5.1 Introduction.....	77
5.2 Experimental Method	78
5.2.1 Apparatus	78
5.2.2 Procedure	79
5.2.3 Manipulation check.....	79
5.2.4 Measures	80
5.3 Results.....	81
5.3.1 Subjective Assessments of Rea.....	81
5.3.2 Effects on Task Measure.....	83
5.3.3 Gender Effects	83
5.3.4 Qualitative Results	83
5.4 Discussion.....	85
5.5 Summary	87
Chapter 6. GENERATING NONVERBAL BEHAVIOR FOR FACE-TO-FACE RELATIONAL DIALOGUE.....	89

6.1 Introduction.....	89
6.2 The Role of Nonverbal Behavior in Face-to-Face Interaction	89
6.2.1. Emotional Display Functions.....	90
6.2.2 Attitudinal Functions	91
6.2.3 Coverbal Interactional Functions.....	92
6.2.4 Coverbal Propositional Functions.....	93
6.2.5 Coverbal Framing Functions	93
6.3 BEAT	94
6.3.1 Knowledge Base	97
6.3.2 Discourse Model.....	98
6.3.3 Language Tagging	98
6.3.4 Behavior Suggestion.....	99
6.3.5 Behavior Selection.....	103
6.3.6 Behavior Scheduling and Animation.....	103
6.3.7 Extensibility	105
6.4 Extensions to BEAT for Relational Agents.....	106
6.4.1 Attitudinal Functions	106
6.4.2 Framing Functions	107
6.5 SUMMARY.....	107
Chapter 7. LONG-TERM RELATIONSHIPS AND HELPING	111
7.1 Introduction.....	111
7.2 Building the Working Alliance.....	112
7.2.1 Scheduling of Techniques.....	114
7.3 Exercise Behavior Change.....	114
7.4 Ethnomethodological Studies of Trainer-Client Interactions	116
7.4.1 Pilot 1: Half-hour face-to-face interactions	117
7.4.2 Pilot 2: Ten-minute text-chat interactions	119
7.5 Summary.....	122
Chapter 8. DESIGN OF A RELATIONAL AGENT FOR LONG-TERM HELPING.....	125
8.1 Requirements and Specifications.....	125
8.2 System Architecture.....	127
8.3 FitTrack ECA Requirements and Specifications.....	129
8.4 Exercise Advisor Agent "Laura"	131
8.4.1 Character Selection Survey.....	133
8.4.2 Nonverbal Behavior Survey.....	134
8.4.3 Facial Expression Survey.....	136
8.4.4 Character Name Survey	137
8.5 Exercise Advisor Server	137
8.5.1 Web Server	137
8.5.2 Dialogue Server	138
8.6 Dialogue Scripting Language and Compiler	140
8.7 Persistence	143
8.8 Exercise Advisor Dialogue Design.....	145
8.9 Summary.....	146

Chapter 9. AN EVALUATION OF LONG-TERM RELATIONSHIPS WITH RELATIONAL AGENTS	151
9.1 Introduction.....	151
9.1.1 Hypotheses.....	152
9.2 Experimental Method	153
9.2.1 Subjects.....	155
9.2.2 Apparatus	158
9.2.3 Procedure	160
9.2.4 Measures	162
9.2.5 Anomalies	164
9.3 Quantitative Results for All Subjects.....	164
9.3.1 Relational Measures for All Subjects	164
9.3.2 Behavioral Outcome Measures for All Subjects.....	169
9.3.3 Participation Measures for All Subjects	171
9.3.4 Exercise Attitudinal Measures for All Subjects.....	171
9.3.5 Other Quantitative Measures for All Subjects.....	173
9.4 Quantitative Results for Sedentary Subjects.....	173
9.4.1 Relational Measures for Sedentary Subjects	175
9.4.2 Behavioral Outcome Measures for Sedentary Subjects.....	188
9.4.3 Participation Measures for Sedentary Subjects	181
9.4.4 Exercise Attitudinal Measures for Sedentary Subjects.....	182
9.4.5 Other Quantitative Measures for Sedentary Subjects.....	182
9.5 Qualitative Analyses	183
9.5.1 Overall Impressions	183
9.5.2 Acceptance of Laura	184
9.5.3 Repetitiveness	185
9.5.4 Caring.....	186
9.5.5 Laura vs. Human Trainer	187
9.5.6 Laura as Motivator.....	188
9.6 Discussion.....	189
9.6.1 Effects on Relationship.....	189
9.6.2 Effects on Physical Activity	190
9.6.3 Repetitiveness	192
9.6.4. Reactions to Laura and FitTrack.....	193
9.6.5 Generality.....	194
9.6.6 Laura version 2.0	194
9.7 Summary.....	195
Chapter 10. CONCLUSION.....	197
10.1 Summary of Significant Contributions.....	197
10.2 Practical Implications	198
10.2.1 Implications for Human-Computer Interface Design	198
10.2.2 Implications for Experimental Design.....	198
10.3 Sociological Implications	199
10.4 Research Implications.....	199
10.4.1 Psychology.....	199

10.4.2 Linguistics.....	199
10.4.3 Artificial Intelligence.....	200
10.5 Ethical Considerations.....	200
10.6 Future Directions.....	201
10.7 Concluding Remarks.....	203
References.....	205
Appendix A. Experimental Protocol for the Small Talk Experiment.....	221
Subject Recruiting Advertisement.....	221
Experimenter’s Script.....	222
Consent Form.....	225
Questionnaire Part 1.....	227
Questionnaire Part 2.....	231
REA Script for SOCIAL Condition.....	234
REA Script for TASK Condition.....	237
Appendix B. Experimental Protocol for the Exercise Advisor Experiment.....	239
Subject Recruiting Advertisements.....	239
Experimenter’s Intake Script.....	242
Consent Form.....	245
Intake Screening Form.....	247
Intake Web Questionnaires.....	248
Subject Take Home Instructions.....	249
Activity Log Web Form.....	252
Activity Log Instructions Web Page.....	253
Self Efficacy Web Form.....	254
Decisional Balance Web Form.....	255
Working Alliance Inventory Web Forms.....	256
Non-Control Evaluation 1 Web Form.....	259
Control Evaluation 1 Web Form.....	260
Non-Control Evaluation 2 Web Form.....	261
Control Evaluation 2 Web Form.....	262
Appendix C. Fragment of Exercise Advisor Scripts.....	263
Script Primitives.....	263
"Int1" Script.....	264
"Int2" Script.....	266
"Int3" Script.....	268
"Int4" Script.....	270
"Int30" Script.....	272
"Farewell" Script.....	273
"FindState" Script.....	274
"FollowUp" Script.....	276
"GetCommitment" Script.....	279
"Greeting" Script.....	284

Chapter 1

INTRODUCTION

People need support from others in order to thrive, but in our society many cannot get the support they need. As our population becomes more mobile and fractionated, and especially as it ages, individuals may find it more and more difficult to get the emotional, social and instrumental support they need to lead healthy, happy and productive lives. Computer agents may be able to fill this growing void, by providing many of the psychological benefits to their users that are known to accrue from having trusting, caring relationships with other people. What kinds of relationships can people have with computers? What can computer agents do to actively draw people into relationships with them? What are the benefits to the people who engage in such relationships? This thesis provides some answers these questions through a series of explorations in building and evaluating such "relational agents."

Relational agents could be used in applications in which a human-computer relationship is established as an end in itself, for example by providing a nurturing outlet (as in Tamagotchis) or partially fulfilling a user's need for intimacy. However, human-human relationships also play an instrumental role in many kinds of tasks, including situations in which a person is attempting to undergo significant cognitive, emotional or behavioral change, such as in education or psychotherapy. This thesis will focus primarily on this second category of applications, in which relationship-building is performed primarily to assist in effecting a desired and beneficial task outcome or change in the user.

1.1 Motivation: The Importance of Personal Relationships

On the premise that people respond to computer agents in fundamentally social ways (following Reeves & Nass (Reeves & Nass, 1996)), a range of applications for relational agents can begin to be delimited by investigating the range of things that human relationships are good for. Provision models of relationships in social psychology give an idea of the possibilities. Some of the types of support that relationships have been found to provide are: emotional support (e.g., esteem, reassurance of worth, affection, attachment, intimacy); appraisal support (e.g., advice and guidance, information, feedback); instrumental support (e.g., material assistance); group belonging; opportunities to nurture; autonomy support; and social network support (e.g., providing introductions to other people) (Berscheid & Reis, 1998). Relational agents could be constructed to provide almost any of these kinds of support.

A large amount of empirical work has been done in social psychology and other fields that demonstrate a significant association between social support and health and survival. In addition to general health and well-being, social support has also been shown to play a significant role in adjustment to specific illnesses, such as cancer and cardiovascular disease. Some of the features of relationships that have been hypothesized to lead to health benefits include: provision of physical and emotional security; establishment of a frame of reference for social reality; normative and informational social influence; and cooperative goal-directed activity. Health and well-being may also be augmented simply because relationships are emotionally gratifying (Berscheid & Reis, 1998). Relational agents could play a significant

role in helping individuals--especially those in acute need (e.g., suffering from an illness and not having any human support network)--cope with their illnesses, and maintain high levels of well-being.

Relationships can also play a role in persuasion. Trustworthiness and likableness of a source of potentially persuasive information play a significant role in Petty and Cacioppo's Elaboration Likelihood Model (Petty & Wegener, 1998). In this theory, if a decision is of low personal importance (a "peripheral route" decision) then source characteristics--such as trustworthiness and likableness of the source of information--have a significant influence on the decision. However, if the outcome of the decision is of high personal importance (a "central route" decision) then these factors have little or no influence on the outcome. Thus, relational agents could be used, for example, as salespeople, which attempt to build relationships with their clients just as good human salespeople do (Anselmi & James E. Zemanek, 1997). Some researchers of personal relationships have also defined interpersonal "closeness" as the degree to which relational partners influence each others' behavior (Kelley, 1983).

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 (Damon & Phelps, 1989). Collaborations between friends involved in these exercises has been shown to provide a more effective learning experience than collaboration between acquaintances (Hartup, 1996). 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.

Even in areas in which the more personal, non-task-oriented, aspects of relationships are downplayed, there is evidence that relationships play an important role in task outcomes. One example of such an area is the world of corporate bureaucracy. Even here, 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 (Gabarro, 1990).

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 customers 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 (Gutek, Cherry, Bhappu, Schneider, & Woolf, 2000). 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.

Finally, although some level of trust is important in all human-computer and human-human interactions (Cassell & Bickmore, 2000), trust and engagement are especially crucial in applications in which a change in the user is desired and which require significant cognitive, emotional or motivational effort on the part of the user. In the helping professions--including clinical psychology, counseling, and coaching--there is a well-documented association between the quality of professional-client relationship and outcomes (Okun, 1997). The positive effect of a good therapist-patient relationship on psychotherapeutic outcomes has been demonstrated in several studies, and has even been hypothesized to be *the* common factor underlying the many diverse approaches to psychotherapy that seem to provide approximately equal results (Gelso & Hayes, 1998). This was even recognized by Freud, who said "It remains the first aim of the treatment to attach him [the patient] to it and to the person of the doctor" (Freud, 1913). Even though different approaches to therapy take different stands on the centrality of the client-therapist relationship (some, such as Interpersonal Therapy see the relationship as not only central to therapy, but relationships in general as central to the patient's sense of self (Kiesler, 1982)), all approaches (including cognitive-behavioral) at least acknowledge that a solid relationship is a pre-requisite for a positive therapeutic outcome. Thus, computer agents that function in helping roles, especially in applications in which the user is attempting to undergo a change in behavior or cognitive or emotional state, could be much more effective if they first attempted to build trusting, empathetic relationships with their users.

A number of instruments have been developed for use in clinical psychotherapy to measure the quality of the client-therapist relationship. One of the most commonly-referred to measures in the literature is the Working Alliance Inventory, which measures the trust and belief that the therapist and patient have in each other as team-members in achieving a desired outcome (Horvath & Greenberg, 1989). This inventory (and similar measures) has been used in therapy to assess the impact of the alliance on problems as wide-ranging as alcoholism, depression, drug use, and personality disorders, and has been demonstrated to have a significant correlation with outcome measures ranging from percentage of days abstinent, drinks per drinking day, and treatment participation (weeks in program) for alcoholism, to employment and compliance with medication, to more general measures such as premature termination, Global Rating Scale (GRS), Global Assessment Scale (GAS), MMPI, Cooley & Lajoy (C-L), and many, many others (Bachelor, 1991; Connors, Carroll, DiClemente, & Longabaugh, 1997; Gaston, 1990; Henry & Strupp, 1994; Horvath, 1994; Horvath & Luborsky, 1993; Horvath & Symonds, 1991; Keijsers, Schaap, & Hoogduin, 2000; Luborsky, 1994; Mallinckrodt, 1003; Raue & Goldfried, 1994).

1.2 Relational Agents

Relational agents are computational artifacts designed to build and maintain long-term, social-emotional relationships with their users. These can be purely software humanoid animated agents--as developed in this work--but they can also be non-humanoid or embodied in various physical forms, from robots, to pets, to jewelry, clothing, hand-held, and other interactive devices. Central to the notion of relationship is that it is a persistent construct, spanning multiple interactions; thus, relational agents are explicitly designed to remember past history and manage future expectations in their interactions with users. Finally, relationships are fundamentally social and emotional, and detailed knowledge of human

social psychology--with a particular emphasis on the role of affect--must be incorporated into these agents if they are to effectively leverage the mechanisms of human social cognition in order to build relationships in the most natural manner possible.

1.3 Embodied Conversational Agents

Although relational agents could be constructed using media ranging from simple text interfaces to speech interaction with autonomous robots, this thesis focuses on relational agents as a specialized kind of embodied conversational agent, which are animated humanoid software agents that use speech, gaze, gesture, intonation and other nonverbal modalities to emulate the experience of human face-to-face conversation. This is motivated by the fact that language is the primary modality used to build human relationships (Duck, 1995), that face-to-face conversation is the primary site of human language use, and that many of the relational strategies that humans use within conversation are nonverbal (Andersen & Guerrero, 1998).

1.4 Two Domains of Analysis

In this thesis I will focus on two domains of analysis with respect to relational agents: the micro-structure of face-to-face relational conversation and the macro-structure of long-term relationships.

As stated above, relationships are primarily negotiated within the context of face-to-face conversation, and relational status permeates all aspects of language use. Thus relational agents must be competent at the encoding and decoding of relational stance, the use of verbal and nonverbal strategies to change relational status, as well as competency at any task being performed—and all of this must be done within the fraction of a second that conversationalists have to plan and produce their utterances in conversation.

At the other extreme of the temporal spectrum, relationships typically last more than a single conversation, and may last an individual's entire lifetime. People use a variety of strategies to maintain relationships over long periods of time and to bridge the times when they are away from their relational partners. Relational agents must also be adept at using these strategies and changing their relational behavior along expected trajectories over time if they are to be accepted by people as long-term sources of social and emotional support.

Two relational agents have been developed to address each of these areas of investigation (see Table 1-1). The REA agent is used as a platform to investigate real-time planning of social dialogue and its effects on users' reported trust in the agent. The Laura agent is used as a platform to investigate long-term relationships, and the impact of relationship maintenance behaviors on users' reported working alliance with the agent.

1.5 Summary

In this chapter I have motivated the development of relational agents, provided a provisional definition for them, and talked about two important levels of analysis that should be taken into account when designing them.



Relational Agent	REA	Laura
Research Focus	Face-to-face conversation	Long-term relationship
Relational Model	Dimensional	Time Common ground
Dialogue Planner	Activation networks	Augmented transition networks
Relational behavior	Small talk	Many (but focus on maintenance)
Task Domain	Real estate	Exercise adoption
Evaluation	Trust	Working alliance
Discussed in	Chapters 4-5	Chapters 6-9

Table 1-1. Two Relational Agents Developed and Evaluated

The remainder of this thesis is broken down into three large parts—background and theory; the design and evaluation of a relational agent for face-to-face interaction (REA); and the design and evaluation of a relational agent for use in a long-term helping relationship (Laura).

- *Chapter 2:* I present previous work in designing artifacts (computational and otherwise) intended to draw users into a sense of relationship with them, and results of studies on the relationships people have with inanimate objects, including computers and computer characters.
- *Chapter 3:* I summarize work in social psychology, philosophy, linguistics and artificial intelligence on the nature of personal relationships, and use this work to derive theoretical models that can be used as the basis for generating relational behavior.
- *Chapter 4:* I present a dialogue planner that can produce natural, mixed task and social dialogue for an embodied relational agent in conversation with a user.
- *Chapter 5:* I present an evaluation of the output of the dialogue planner from Chapter 4. This human subjects study investigates the effects of social dialogue on subjects' trust in and perception of a life-sized embodied conversational agent.
- *Chapter 6:* I present a software architecture for producing natural, conversational nonverbal behavior for an embodied relational agent.

- *Chapter 7:* I present motivation for developing relational agents for long-term health behavior change applications—with a particular emphasis on exercise adoption—and the results of two studies of exercise trainers and their interactions with clients.
- *Chapter 8:* I present the design of a relational agent that is able to maintain a long-term relationship with users, in the domain of exercise adoption.
- *Chapter 9:* I present an evaluation of the relational agent from Chapter 8, involving 100 subjects who interact with the agent daily for a month.
- *Chapter 10:* I summarize my results, discuss implications of these results for a number of disciplines, and present future work for relational agents.

Chapter 2

RELATED WORK

People claim to have relationships with their pets, their cars, their computers and many other non-human entities, whether animate or inanimate. Do they really? And, what does it mean if they do? This chapter reviews previous studies and analyses of relationships people have with non-humans, with a special focus on studies of people's relationships with computers and computer characters.

2.1 Introduction

Before reviewing work on different aspects and types of relationship it helps to have a clear definition of what a relationship is. Based on work in the social psychology of personal relationships (discussed further in Chapter 3), a provisional definition of relationship is that it is something that describes a unique pattern of interdependent behavior between two entities, whether they be people or not, animate or not. A key aspect of relational agents is that they are artifacts intentionally designed to draw people into a sense of relationship with them. Thus, in the following sections I will present related research and development work broken down into categories of non-relational artifacts (things that people may feel a sense of relationship with, but were not intentionally designed with that purpose in mind), and artifacts which have varying degrees of relational agency, with the latter category broken down into social artifacts (intentionally designed to display social cues or engage people's social cognition), affective agents (intentionally designed to display or recognize affective cues or to manage a user's emotional state), anthropomorphic agents (designed to have a human physical form), and truly relational agents. These categories are not mutually exclusive and there are many artifacts that span several, if not all of these categories. I conclude with discussion of related work that does not fit in any of the above categories, as well as some ethnographic studies of people's attitudes towards these kinds of agents and how they have changed over time.

2.2 Relationships with Non-relational Entities

People claim to have relationships with their cars, their tools, and other non-human or inanimate entities. For the most part, these are not true relational agents (according to the definition above) since they were not intentionally designed to produce that response in the people who interact with them. Exactly what do people mean when they say this?

Csikszentmihalyi and Halton's book on "The Meaning of Things" gives perhaps the most comprehensive treatment to date on the psychology of interactions between people and man-made objects (Csikszentmihalyi & Rochberg-Halton, 1998). According to them, much of the work in psychology on the nature of people's interactions with objects is mostly concerned with objects as symbolic representations for the self, for others or for relationships (e.g., Freud, Jung, and even Winnicott's treatment of "transitional objects" (Winnicott, 1982)), but are not at all concerned with the actual experience that people have with concrete objects in

the world. These experiences, though, have a very significant impact on our lives and our psyches:

...man-made objects have an extremely important role to play in human affairs. It is quite obvious that interaction with objects alters the pattern of life; for instance, that refrigerators have revolutionized shopping and eating habits, that automobiles created suburbs and increased geographical mobility, or that television is changing how family members relate to one another. (Csikszentmihalyi & Rochberg-Halton, 1998, pg. 14)

Each new object changes the way people organize and experience their lives. (ibid, pg. 46)

Looking back at the provisional definition of relationship as “a unique pattern of interdependent behavior”, it can be seen that our interaction with the objects around us can change our behavior and, to the extent that these objects can be said to have behavior, this behavior can be dependent upon what we do with them. Thus, the relational criteria of interdependent behavior can certainly be satisfied. As for uniqueness, there are objects in all of our lives that are, to varying degrees, unique:

My old living room chair with its worn velvet fabric, musty smell, creaking springs, and warm support has often shaped signs in my awareness. These signs are part of what organizes my consciousness, and because my self is inseparable from the sign process that constitutes consciousness, that chair is as much a part of my self as anything can possibly be. (ibid, pg. 14)

To the extent, then, that our behavior is dependent upon an object, and that pattern is unique, we can say we have a relationship with that object. To the extent that these criteria are satisfied, we can also say that the object “has meaning” to us:

When a thing “means something” to someone, it is interpreted in the context of past experiences, either consciously, or unconsciously in the form of habit. (ibid, pg. 21)

Pets tend to satisfy both of these criteria quite well, which may explain their popularity and the value many people place in their relationships with them. Objects that do not meet these criteria include non-unique objects (e.g., any disposable commodity), objects that we have not established a pattern of interaction with, or objects that do not influence our behavior in any significant way.

2.2.1 Trust in Man-Made Artifacts

Trust is one of the most important qualities of any relationship (elaborated further in Chapter 3). There has been a fair amount of work over the last few decades on people’s perceptions of trust in man-made artifacts, particularly in machinery and, more recently, computers. Tseng and Fogg define trust as “a positive belief about the perceived reliability of, dependability of, and confidence in a person, object, or process,” and claim that it is one of the key components used in assessments of “computer credibility” (Tseng & Fogg, 1999).

Research on human-computer interfaces has found several interesting results with respect to trust. It has been found that trust in intelligent systems is higher for systems that can explain and justify their decisions (Miller & Larson, 1992). There have also been studies showing how specific design elements, such as the use of color and clipart (Kim & Moon, 1997) or the inclusion of comprehensive product information (Lee, Kim, & Moon, 2000) can influence a user’s perception of trust in an interface. In anthropomorphic interfaces, pedagogical agents, especially those that are highly expressive, have been found to affect students’ perceptions of trust; such agents are perceived as helpful, believable, and concerned

(Lester, Converse et al., 1997). However, Mulken, et al, found that personification of an interface by itself does not appear to be a sufficient condition for raising the trustworthiness of a computer (Mulken, Andre, & Muller, 1999).

2.3 Social Agents

Here, I define “social agents” as those artifacts, primarily computational, that are intentionally designed to display social cues or otherwise to produce a social response in the person using them.

The seminal work in this area was a series of studies by Reeves and Nass who demonstrated that people respond in social ways to computers (and other media) when provided with the appropriate social cues, even though they are typically unconscious of this behavior (Reeves & Nass, 1996). Examples of their studies are investigations into politeness behavior, proximity effects, and gender effects (some additional studies are described in section 2.5 below). In the politeness experiment, they demonstrated that people were more polite when giving evaluations to a computer about its performance than when giving the same evaluation via a different computer, as would be predicted from people’s behavior when evaluating other people. In the proximity experiment, they demonstrated that people respond to the perceived distance between themselves and someone in a picture in the same way they would respond to actual distance between themselves and another person; their evaluations are more intense, they pay more attention, and they remembered the observed person better when they appear closer. In the gender effects study they demonstrated that people reacted with biases congruent with gender stereotypes when evaluating computers with male vs. female voices. The research paradigm followed in all of this work was to take a study from the field of social psychology, replace one of the people in the study with a computer, and repeat the study to show that the other people would behave in the same way they would have if they had been interacting with a person. This body of work was used as the motivation for several commercial products (e.g., Microsoft “Bob” and the Microsoft Office Assistant) and inspired the emergence of the entire field of “social computing.”

There have been numerous workshops over the last few years on the topics of “social agents,” “social cognition,” “social adeptness,” and “computer etiquette,” demonstrating an increased interest in not only studying social responses in human-computer interaction, but in actually building artifacts that provoke these responses. Some recent examples of work in this area include conveyance of personality and impression management through agent gaze behavior (Fukayama, Ohno, Mukawa, Sawaki, & Hagita, 2002), and the use of familiar social protocols for meeting management (Yan & Selker, 2000).

2.4 Affective Agents

Affective agents are those intentionally designed to display affect, recognize affect in users, or manipulate the user’s affective state. Thus, these fall entirely into what Picard terms “affective computing”; computing that relates to, arises from, or deliberately influences emotions (Picard, 1997). Affect and “emotional intelligence” (Salovey & Mayer, 1990), play crucial roles in relationships, and are thus important components of relational agents.

Several researchers have explored the development of technologies for sensing user affect through a variety of physiological, nonverbal and paraverbal channels, including: facial expression (Kapoor & Picard, 2002); posture (Kapoor, Mota, & Picard, 2001); galvanic skin

response (Picard & Scheirer, 2001); muscle contraction (e.g., as measured through grip pressure on a computer mouse (Reynolds & Picard, 2001)); and speech (Scherer, 1981; Scherer, Ladd, & Silverman, 1984; Williams & Stevens, 1972); and some have worked on techniques for fusing information from a number of these modalities (Picard, Vyzas, & Healey, 2001). Some have also investigated models for inferring affective tone from natural language text (Liu, Lieberman, & Selker, 2003).

Other researchers have developed systems for displaying affective signals using a variety of modalities, including: speech (Cahn, 1990); facial expression (Pelachaud, Badler, & Steedman, 1994); motion dynamics (Rose, Bodenheimer, & Cohen, 1998); and natural language text (Hovy, 1986).

There have been several conversational systems developed that attempt to convey emotion in Embodied Conversational Agents (discussed in detail below) via facial expression and/or body posture. Examples are the Cosmo pedagogical agent (Lester, Towns, Callaway, Voerman, & Fitzgerald, 2000), the Byrne robocup sportcaster (Binstead, 1998), and others (Andre, Muller, & Rist, 1996; Beskow & McGlashan, 1997).

Some work has also been done to integrate affect sensing and production. Ball and Breese describe a sophisticated system for recognizing user affect and personality, as well as generating affect and personality using a variety of behavioral cues and a single Bayesian belief network to perform both recognition and production (Ball & Breese, 2000). Their cues currently include vocal cues (average pitch, pitch range, speech speed, speech energy), verbal cues (active, positive, strong, terse, or formal aspects of lexical choice), facial expression, gesture (speed and size) and postural information. The affective variables currently modeled are valence and arousal, and the personality variables are friendliness and dominance/submissiveness.

Systems for managing user affective state have the greatest relevance for the current thesis, unfortunately there are few examples of these in the literature. The best work in this area is the CASPER affect-management agent developed by Klein (Klein, Moon, & Picard, 2002; Klein, 1999), which was demonstrated to provide relief to users experiencing frustration. The system presented a frustrated user with a series of menus that prompted the user to describe their affective state, provided paraphrased feedback, allowed users to repair the computer's assessment and provided empathetic and sympathetic feedback. This agent was found to be significantly better than a venting-only agent (to which users could simply describe how they felt in an open-ended manner without feedback), or an agent that ignored their emotions completely, in relieving frustration, as measured by the length of time users were willing to continue working with a computer after a frustrating experience.

2.5 Anthropomorphic Agents

Anthropomorphic agents are systems intentionally designed to have human form. They overlap with the previous two categories since the very existence of a body sends a strong social cue to the user about what to expect in the interaction, and since users will infer personality and affective state from any embodied agent, regardless of how it behaves.

Several studies have been done to determine if there is a "persona effect", that is, if the mere presence of a face or body in the interface has a significant impact on user attitudes or behavior. Koda and Maes (Koda & Maes, 1996) and Takeuchi and Naito (Takeuchi & Naito, 1995) studied interfaces with static or animated faces, and found that users rated them to be

more engaging and entertaining than functionally equivalent interfaces without a face. Kiesler and Sproull (Kiesler & Sproull, 1997) found that users were more likely to be cooperative with an interface agent when it had a human face (vs. a dog or cartoon dog).

In terms of social behaviors, Sproull et al. (Sproull, Subramani, Kiesler, Walker, & Waters, 1997) showed that subjects rated a female embodied interface significantly lower in sociability and gave it a significantly more negative social evaluation compared to a text-only interface. Subjects also reported being less relaxed and assured when interacting with the embodied interface than when interacting with the text interface. Finally, they gave themselves significantly higher scores on social desirability scales, but disclosed less (wrote significantly less and skipped more questions in response to queries by the interface) when interacting with an embodied interface vs. a text-only interface. Men were found to disclose more in the embodied condition and women disclosed more in the text-only condition. Possible explanations for these findings are that the embodied interface was very unfriendly and cold in appearance, leading to the low ratings of sociability and negative social evaluation (even the “pleasant” face used in the experiment looked unfriendly, the face was completely static in between utterances, and a DECTalk text-to-speech synthesizer was used for the agent’s voice, all of which could lead to ratings of unfriendliness). The increased social presence of the embodied interface (relative to text) could have led to subjects feeling less relaxed and assured, and led increases in social desirability effects.

In their survey of user studies on embodied agents, Dehn and van Mulken conclude that there is no "persona effect", that is a general advantage of an interface with an animated agent over one without an animated agent (Dehn & Mulken, 2000). However, they believe that lack of evidence and inconsistencies in the studies performed to date may be attributable to methodological shortcomings and variations in the kinds of animations used, the kinds of comparisons made (control conditions), the specific measures used for the dependent variables, and the task and context of the interaction.

2.5.1 Embodied Conversational Agents

Embodied Conversational Agents (ECAs) are animated humanoid software agents that use speech, gaze, gesture, intonation and other nonverbal modalities to emulate the experience of human face-to-face conversation with their users.

Work on the development of ECAs, as a distinct field of development, is best summarized in (Cassell, Sullivan, Prevost, & Churchill, 2000). In addition to REA (Cassell et al., 1999) (described below), some of the other major ECA systems developed to date are Steve (Rickel & Johnson, 1998), the DFKI Persona (Andre et al., 1996), Olga (Beskow & McGlashan, 1997), Gandalf (Thorisson, 1997), and pedagogical agents developed by Lester, et al, (Lester, Stone, & Stelling, 1999; Lester, Voerman, Towns, & Callaway, 1997). There are also a growing number of commercial ECAs, such as those developed by Extempo, Headpedal, and Artificial Life, and the Ananova newscaster developed by Ananova, Ltd. These systems vary greatly in their linguistic capabilities, input modalities (most are mouse/text/speech input only), and task domains, but all share the common feature that they attempt to engage the user in natural, full-bodied (in some sense) conversation. Microsoft has also produced a toolkit for developing animated talking agents (Microsoft Agent), although these characters are unable to use speech and nonverbal modalities at the same time, making them unusable for natural, multi-modal conversation.

REA--a project I led in the Gesture & Narrative Language Group at the MIT Media Lab from 1999 to 2002--is a real-time, multi-modal, life-sized ECA who plays the role of a real estate agent who can interview potential home buyers and show them around houses she has for sale (see Figure 2-1) (Cassell et al., 1999; Cassell, Bickmore, Vilhjálmsón, & Yan, 2000). Other ECAs developed at the Media Lab include the Sam, GrandChair, and MACK systems. Sam is a peer embodied conversational storyteller for children, which shares a real castle play space and a set of story-evoking toys with a child(Cassell, Ananny et al., 2000). GrandChair uses an attentive child embodied agent listener to elicit and save grandparents' reminiscences (Smith, 2000). The Media Lab Autonomous Conversational Kiosk (MACK) is a virtual receptionist for the Media Lab that can give directions and talk about the research that different groups are doing (Cassell et al., 2002).

There have been numerous studies on the efficacy of ECAs, most comparing similar applications with and without an ECA. Andre, Rist and Muller found that users rated their animated presentation agent ("PPP Persona") as more entertaining and helpful than an equivalent interface without the agent (Andre, Rist, & Muller, 1998). However, there was no difference in actual performance (comprehension and recall of presented material) in interfaces with the agent vs. interfaces without it.

In a user study of the Gandalf system (Cassell, 1999), users rated the smoothness of the interaction and the agent's language skills significantly higher under test conditions in which Gandalf utilized limited conversational behavior (gaze, turn-taking and beat gesture) than when these behaviors were disabled.

2.5.2 Sociable Robots

There has been growing interest in recent years in building anthropomorphic robots, with the motivation being that they are easier for people to interact with when they are in a familiar form.

From a social skills perspective, the most sophisticated humanoid robot constructed to date is Kismet (Breazeal, 2000). Kismet is a humanoid robot designed to engage in socially



Figure 2-1. User Interacting with Rea

situated learning with untrained humans. In order to support this, Kismet models the interactions between mothers and their infants, is able to sense where the user and certain objects are in its environment and the affective tone of the user's voice, and reacts using facial expressions, proximity (moving its head forward or back) and "proto-speech" (meaningless utterances, but with appropriate affective intonation). Kismet has sophisticated emotion, motivation and behavioral systems that allow it to carry on a highly engaging interaction with a user, even without the exchange of propositional content in the speech channel.

Sidner has developed a robotic ECA that performs hosting activities, with a special emphasis on "engagement"--an interactional behavior whose purpose is to establish and maintain the connection between interlocutors during a conversation (Sidner, 2002). The robot—named "MEL"—is driven by a model of collaborative behavior based on studies of human dialogue (Rich & Sidner, 1998).

2.6 Relational Agents

As stated above, relational agents are those intended to produce relational cues or otherwise produce a relational response in their users, such as increased liking for or trust in the agent. Studies by Reeves and Nass and their students on relational aspects of human-computer interaction constitute the bulk of work in this area to date. The majority of these studies use non-embodied, text-only human-computer interfaces.

In their book on the Media Equation, Reeves and Nass demonstrated the following relational effects (Reeves & Nass, 1996):

- Computers that use flattery, or which praise rather than criticize their users are better liked.
- Computers that praise other computers are better liked than computers that praise themselves, and computers that criticize other computers are liked less than computers that criticize themselves.
- Users prefer computers that match them in personality over those that do not (the "similarity attraction" principle).
- Users prefer computers that become more like them over time over those which maintain a consistent level of similarity, even when the resultant similarity is the same.
- Users who are "teamed" with an computer will think better of the computer and cooperate more with it than those who are not teamed (the "in-group membership" effect, which can be achieved by simply signifying that the user and computer are part of a team).

Since the Media Equation was published, Reeves and Nass and their students have continued doing studies within this "Computers As Social Actors" paradigm. Morkes, Kernal and Nass demonstrated that computer agents that use humor are rated as more likable, competent and cooperative than those that do not (Morkes, Kernal, & Nass, 1998). 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 (Moon, 1998).

A number of commercial products—mostly toys—have been developed over the last few years which fall into the category of relational agents in that they are designed to cultivate a

sense of relationship with their users. Most of these artifacts play on people's need to express nurturance by requiring caretaking in order to thrive, or by engaging in familiar social interaction patterns. Many of these artifacts also change their behavior over time or otherwise provide a highly variable, rich set of expressions to give the sense of uniqueness crucial for relationships. Examples include the Tamagotchi (one of the first and simplest, yet wildly successful in Japan), Hasbro's Furby, Sony's AIBO (robotic dog) and iRobot's My Real Baby (robotic baby doll). One of the more interesting entries in this area is Matsushita's Tama, a robotic cat designed to be a conversation partner for elderly people (Drexler, 1999). According to a manager overseeing the project, "The idea [behind Tama] is animal therapy; a network system will enable the pets to speak to the elderly in a natural way, especially to people who are living alone, and this will make them more comfortable." The developers claim that past interactions are recorded in memory, but there is no indication of what this memory is used for.

2.7 Other Related Work

Additional work has been done on modeling aspects of relational interaction, but without the intent or ability of the system to build a relationship between itself and a user.

Pautler developed a model of social perlocutions--the psychosocial effects of various social speech acts--and applied it in the LetterGen system for helping people write socially appropriate email messages (Pautler, 1998). His model included effects on the hearer's emotional state and the effects of these changes on the relationship between the speaker and hearer.

Walker, Cahn and Whittaker developed a system that took a speech act along with the social distance and power relationship between two interactants as inputs, and generated a surface form utterance of the speech act using the appropriate politeness strategy (Walker, Cahn, & Whittaker, 1997).

Elliott developed a multi-agent system in which agents could learn models of the "goals, standards and preferences" of other agents and use these to predict and reason about their emotions (Elliott, 1992).

Colby's PARRY system—developed as a simulation of a schizophrenic paranoid personality—responded to perceived malevolence on the part of the user (in typed text interactions) by responding with fear, anger or mistrust resulting in counterattack or withdrawal (Colby, 1981).

Finally, Affect Control Theory provides a set of empirically-grounded mathematical models for explaining and predicting expressive aspects of action, including how people evaluate the affective tones of different social situations, their identities in those situations, and actions taken by them or others (Heise, 1992).

2.8 Changing Attitudes Towards Relational Agents

Turkle's work over the last twenty years on the effects of computational objects on human developmental psychology is of great relevance to this thesis, since it describes how people have come to accept increasingly complex—and relational—computational artifacts into their lives.

In her book "Life on the Screen", she 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 in general (Turkle, 1995). Reporting on users' experiences with DEPRESSION 2.0 (a program that presents itself as a psychotherapist specializing in treating depression), she says that:

The testimony of people who have used it suggests that although they find it "clunky," they are, in the main, happy to accept its help, no (philosophical) questions asked. People can sit for hours at a stretch talking to this program about the intimate details of their lives and then dismiss the experience, saying, "it's just a machine." (Turkle, 199, pg. 102)

In more recent work she has studied people's acceptance of "relational artifacts," defined as "computational objects designed to recognize and respond to the affective states of human beings—and indeed, to present themselves as having 'affective' states of their own." (Turkle, 2002). These artifacts include toys such as Tamagotchis, Furbies AIBOs and robotic dolls that interact with people on a relational and psychological level, "pushing our evolutionary buttons" by synthesizing emotional displays and social behavior. She finds that users of these systems are less concerned with their internal mechanisms than they are with learning strategies for getting them into desired emotional states.

2.9 Summary

A great deal of work has been done over the last two decades on technologies relevant to relational agents. However, with the exception of some work on commercial products, no one has undertaken the development of an agent intended to build and maintain a relationship with its user over an extended period of time, nor has anyone (with the possible exception of Mitsushita's Tama) investigated the role these relationships could play in effecting instrumental task outcomes.

Chapter 3

THEORY

Although relational agents could be constructed without reliance on a deep theory of what a social-emotional relationship is, such a theory can provide a basis for the necessary flexibility and generality required for an agent that can function adaptively over a wide range of contexts, including multiple interactions over long time spans. In addition, a theory of human-computer relationships can provide a framework within which the design and evaluation of relational agents can be better understood. This chapter integrates work from social psychology, linguistics, communication and artificial intelligence to provide such a theoretical framework.

I will start by reviewing fundamental definitions of personal relationship from the field of social psychology, then discuss artificial intelligence theories of multi-agent collaboration and accommodation that will be used as the basis for a theoretical framework before presenting the synthesized theory of what a relationship is. I then proceed to discuss how this theoretical framework can be applied to understanding aspects of relational behavior in two realms: the micro-structure of face-to-face conversation and the macro-structure of relational maintenance in long-term relationships.

3.1 The Social Psychology of Personal Relationships

Most recent work in the social psychology of personal relationships takes a fundamentally dyadic approach to the concept of “relationship” (Berscheid & Reis, 1998). Kelley *et al* define this concept as referring to two people whose behavior is interdependent, in that a change in the state of one will produce a change in the state of the other (Kelley, 1983). Thus, a relationship does not reside in either partner alone, but in their interaction with each other. Further, a relationship is not defined by generic patterns of interaction (e.g., associated with stereotypical roles), but by the unique patterns of interaction for a particular dyad (Berscheid & Reis, 1998).

In addition to defining relationships in terms of what people actually do together and their degree of interdependence, relationships are also often defined in terms of what the people in them provide for one another. Duck, for example, defines the following list of provisions that “friends” in our culture are expected to provide for each other (Duck, 1991):

- Belonging and a sense of “reliable alliance”. The existence of a bond that can be trusted to be there for a partner when they need it.
- Emotional integration and stability. Friendships provide necessary anchor points for opinions, beliefs and emotional responses.
- Opportunities for each partner to talk about themselves. Friendships help fulfill the need for self-expression and self-disclosure.
- Provision of physical, psychological and emotional support. Physical support involves doing favors, such as giving someone a ride or washing the dishes. Psychological support involves showing appreciation for the other and letting them

know their opinions are valued. Emotional support includes affection, attachment and intimacy.

- Reassurance of worth and value, and an opportunity to help others. We value friends because of their contribution to our self-evaluation and self-esteem, directly via compliments and indirectly by telling us of the good opinions of others. Also, friends increase our self-esteem by simply attending to us, by listening, asking our advice and valuing our opinions.

Similarly, Brehm defines “intimate relationships” in terms of provisions such as intimacy, nurturance, assistance, social integration, and reassurance (Brehm, 1992).

This notion of provisions is also reflected in economic models of relationship, such as exchange theory (Brehm, 1992), to the extent that they model relationships in terms of costs vs. benefits. Social exchange models are economic models of the costs, benefits, investments, and alternatives individuals have in relationships and how these relate to their levels of commitment. Social exchange models have received more empirical validation than any other theoretical framework in the social psychology of personal relationships. In these models the benefits of a relationship can be seen as the perceived value of the provisions that one partner receives while their costs can be seen as the perceived cost of the provisions that they provide to the other. In these theories a relationship only exists when there is such an economic exchange, and the longevity of the relationship (commitment to continue) can be reliably predicted from equations involving these costs and benefits and other terms such as perceived alternatives to and amount of investment in the relationship.

Many researchers have also attempted to define dimensional models that identify the underlying features that characterize different stereotypical relationships. The most commonly mentioned dimensions are power and social distance, with social distance further refined by many researchers into as many as 14 sub-dimensions (Brown & Gilman, 1972; Burgoon & Hale, 1984; Spencer-Oatey, 1996; Svennevig, 1999). Other dimensions used to characterize relationships include equal vs. unequal, hostile vs. friendly, superficial vs. intense, informal vs. formal (Wish, Deutsch, & Kaplan, 1976), and various typologies of love (Brehm, 1992).

Some work has also been done on integrating these various models. For example McGuire derived a set of 72 types of helping behavior and grouped them into four factors—casual helping (low cost to the helper), substantial personal helping (substantial cost to the helper), emotional helping, and emergency helping—and then determined how the frequency of these varied with the interpersonal closeness of the relationship between the helper and helpee, effectively bridging provisional and dimensional models of relationship (McGuire, 1994).

3.2 Multi-Agent Collaboration

Relationships involve collaborative behavior; coordinated activity in which the participants work jointly with each other to satisfy a shared goal. Examples of such 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). To begin to formalize a notion of relationship for relational agents, I now turn to the field of artificial intelligence and formalisms for planning and collaboration among autonomous agents.

Perhaps the most complete theory of multi-agent collaboration in the field of AI is the SharedPlans theory developed by Grosz & Sidner, et al. (Grosz & Kraus, 1993; B. Grosz & S. Kraus, 1996; Grosz & Sidner, 1990). This work describes a normative, performance model detailing how agents move from individual goals and intentions into collaborative, coordinated activity, based on representations in the minds of the individual agents.

The SharedPlans theory is based on the theories of Bratman and Pollack, who outline a mental-state view of plans in which having a plan is not just knowing how to do an action, but also having the intention to do the actions entailed (Bratman, 1990; Pollack, 1990). This approach thus differentiates between knowing *how* to accomplish a goal (a "recipe") and having a plan, which includes intentions. The SharedPlans theory explains how a group of agents can incrementally form and execute a SharedPlan that then guides and coordinates their activity towards the accomplishment of a shared goal.

Informally, two agents are said to have a SharedPlan when they mutually believe that: a) they have a common goal; b) they have agreed on a recipe to accomplish the goal; c) they are each capable of performing their assigned actions; d) each intends to do their assigned actions; and e) they are committed to the overall success of the collaboration. SharedPlans are usually incrementally refined and executed by the collaborating agents; in a typical scenario a SharedPlan is initially partial (incompletely specified) and only becomes completely specified once the agents have finished refining and executing it. Refinement of a partial plan is carried out through means-ends reasoning and negotiation among the agents.

The formalization presented in (B. J. Grosz & S. Kraus, 1996) uses a first-order logic augmented with several modal operators, meta-predicates, and action expressions. Table 3-1 lists the notations from that work that will be used here (I am abstracting from their representation; reference to the time, context and plan identifier parameters are not needed for the current discussion and have been omitted). In this representation, actions can be readily executable ("basic-level") or complex, with complex actions having recipes consisting of additional sub-actions that are basic-level or complex. Thus actions and their decomposition form recipe trees, which reflect hierarchical plan decomposition.

The theories rely on two different notions of intention. Both notions follow Bratman in that they prevent the agent(s) from adopting conflicting intentions. "Intend to" (Int.To) is an attitude a single agent holds with respect to an action, requires that the agent know a way of doing the action (e.g., a recipe) and commits it to means-end reasoning, if required, to carry the action out. An agent cannot intend for another agent to perform an action using this attitude. "Intend that" (Int.Th) is an attitude one or more agents hold with respect to a proposition, and reflects the desire for the agent(s) to achieve the state of the world specified, however, it is not as strong a commitment as intend-to, and does not require that the agent(s) know how to achieve the action and does not commit them to means-end reasoning. Further, an agent can intend-that another agent achieve the specified proposition.

Type	Notation	Meaning
Literals	G1, G2	Agents
	GR	a group (set) of agents
	P	a state of the world (a proposition)
	α	an action
	R_α	a recipe to achieve α
Modal Operators	Bel(G,P)	Belief
	MB(GR,P)	mutual belief
	Int.To(G, α)	Agent intends to perform α .
	Int.Th(GR,P)	Agent(s) intend that state P be achieved.
	Do(GR, α)	Agent(s) G do α .
Meta-predicates	FSP(GR, α , R_α)	Agents have a full shared plan to do α using R_α .
	PSP(GR, α)	Agents have a partial shared plan to do α .
	SP(GR, α)	The group has a FSP or a PSP and a FSP to elaborate it.
	CBAG(GR, α , R_α)	A group of agents “can bring about” α using Ra.

Table 3-1. SharedPlans Notations Used (B. J. Grosz & S. Kraus, 1996)

Grosz and Kraus define a Full Shared Plan as a meta-predicate that represents the situation in which two (or more) agents have determined the recipe by which they are going to achieve some action, and the members of the group have adopted intention-to toward all of the basic-level actions in the recipe as well as intentions-that toward the actions of the group and its other members. A Full Shared Plan is represented as $FSP(GR, \alpha, R_\alpha)$, where: GR is the group of agents involved in the collaboration; α is the action to be achieved, and R_α is the recipe to be used. A Partial Shared Plan— $PSP(GR, \alpha)$ —differs from a Full Shared Plan in that the agents may only have a partial recipe for doing the action, they may have only partial plans for doing some of the steps in the recipe, and there may be some sub-actions which the have not yet been assigned to a particular agent.

While relationships certainly involve collaborative behavior, it is worth pointing out at this point that a collaboration, as modeled by a SharedPlan, is not a good representation for the concept of “having a personal relationship”. Although not fully specified here, SharedPlans involve 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 (e.g., as provisions) 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).

Grosz and Sidner discuss general strategies for moving from one agent’s having a goal to a group of agents having a SharedPlan to achieve that goal (Grosz & Sidner, 1990). They bridge part of this gap via a “conversational default rule” (a rule which operates in the

absence of evidence to the contrary) which states that if one of the agents has a goal for P and they are cooperative and communicating about this desire to achieve P, then they will come to mutually believe that this agent has a desire the group to construct a SharedPlan to achieve P. However, to move from this mutual belief about G1's desire for a collaboration to actually performing the collaboration requires (probable) negotiation and assent by G2. How can agents assess whether their potential collaborators (such as G2) will be likely to provide assistance? These expectations of future collaborations can be defined by appeal to the notion of accommodation, which I turn to next.

3.3 Accommodation

Accommodation is 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. It is hypothesized to underlie a wide range of collaborative and linguistic behavior. Thomason defines accommodation as follows:

Most generally, accommodation consists in acting to remove obstacles to the achievement of desires or goals that we attribute to others... I am accommodating you, for instance, if I open the door when I see you approach it with your hands full of packages. (Thomason, 1990)

This notion can be directly encoded in the SharedPlans formalism above as the following meta-predicate, representing the situation when G2 accommodates goal P for G1:

$$\begin{aligned} \text{Acc1}(G1,G2,P) &\Leftrightarrow \\ (1) \quad &\text{Bel}(G2,\text{Int.Th}(G1,P)) \ \& \ \neg\text{Bel}(G2,P) \ \& \\ (2) \quad &\text{Bel}(G2,(\exists R_\alpha)[[\text{Do}(\{G1,G2\},\alpha)\Rightarrow P] \ \& \ \text{CBAG}(\{G1,G2\},\alpha,R_\alpha)]) \\ (3) \quad &\Rightarrow \text{Int.Th}(G2,\text{SP}(\{G1,G2\},\alpha)) \end{aligned}$$

That is, whenever (1) G2 believes that G1 desires to achieve P (and believes that P is not already satisfied) and (2) believes it is able to help (believes there is a recipe whose execution will result in P and which the agents are able to collaborate on in order to achieve P) then (3) G2 will adopt the intention to collaborate with G1 on P's achievement.

Using this notion of accommodation, while a good starting place, has several shortcomings if it is to be used as the basis for a provision-based definition of relationship. First, in order for agents to be able to 'keep score' of when other agents have actually helped them in the past, accommodation needs to be based on mutual knowledge, rather than just on beliefs in G2's head (i.e., G2 doesn't get 'credit' for helping unless everyone is aware that helping is going on). Second, relationships are not only defined by the goals each partner will help the other with, but the manner in which these goals are satisfied; namely by the set of recipes that the partners will use to help the other. That is, G2 may not be willing to use any recipe at its disposal in helping G1 with P, but only a specific subset of those it knows about. Further, it may be characteristic of this dyad for them to use a particular recipe together when performing a particular kind of task.

A second definition of accommodation that captures these additional subtleties is:

- $$\text{Acc2}(G1, G2, P, R_\alpha) \Leftrightarrow$$
- (1) $\text{MB}(\{G1, G2\}, (\exists R_\alpha)[\text{Do}(\{G1, G2\}, \alpha) \Rightarrow P] \ \&$
 - (2) $\text{MB}(\text{Int.Th}(G1, P) \ \& \ \neg \text{Bel}(G1, P) \ \& \ \neg \text{Bel}(G2, P) \ \&$
 - (3) $\text{CBAG}(\{G1, G2\}, \alpha, R_\alpha)]$
 - (4) $\Rightarrow \text{Int.Th}(G2, \text{SP}(\{G1, G2\}, \alpha))$

In this definition, (1) both agents must believe there is a specific recipe that will bring about P, and (2) when there is mutual belief that G1 desires P (and P is not already satisfied), and (3) they can achieve P by collaborating on the execution of the recipe, then (4) G2 will adopt the intention to collaborate with G1 on P's achievement. Note that the condition in (3) lets G2 "off the hook" for accommodating if it is presently unable to help out.

Essentially, this is an expectation, a future potential for action. The inner MB indicates that it's not enough for G2 to simply observe that G1 desires P (e.g., through observation or third parties) and act in secrecy to aid G2. In relational accommodation both agents need to be aware that accommodation is happening. As mentioned above, both agents are very interested in 'keeping score' of when others actually accommodate them (an important part of relationships, for example, as represented in exchange models), and this can't happen unless everyone is aware that accommodation is intentionally going on. The outer MB allows the agents to plan their lives in such a way that they can rely on the accommodation of others. It turns out that line (4) is actually too strong, but it's the ideal case. In reality, there is a great deal of uncertainty about what one agent will help another one with, and agents constantly assess the strength of these accommodation relations.

3.4 A Formal Definition of Relationship

Given the definition of relational accommodation above, we can proceed to define a relationship between two agents as the set of all such accommodation relations that hold between them:

$$\text{Relationship}(G1, G2) \Leftrightarrow \{ \langle x, y, P, R_\alpha \rangle \mid x, y \text{ in } \{G1, G2\} \ \& \ R_\alpha \in \text{Recipes} \ \& \ \text{Acc2}(x, y, P, R_\alpha) \}$$

Two agents are said to 'have a relationship' if this set is non-empty.

Under this definition, "having a relationship" or "being in a relationship" is not a collaboration (SharedPlan), but is a set of expectations for future collaborations should the need arise. On the other hand, "establishing a relationship", "changing a relationship", "maintaining a relationship" are specific goals which require collaboration (and SharedPlans) in order to be accomplished, at the time they are performed.

As mentioned above, the normal state of most human relationships (excluding formalized ones such as contractual relationships) is that there is a great deal of uncertainty regarding which accommodation relations actually hold. That is, an agent is constantly trying to update its degree of belief in each of the things that its partner will accommodate it on, as well as its own likelihood of accommodating its partner on a variety of activities. It is these assessments, and the desire for information about them, that drives most relational behavior in human social interaction. And, it is the agreement on these assessments between partners and their

degree of overlap with stereotypical relationships that ultimately defines the relationship between them.

Following the observations from the social psychology of personal relationships in section 3.1, the particular activities that two individuals expect to engage in with each other are (potentially) unique to that dyad, however to the extent that this set of activities overlaps with a common relational stereotype (e.g., "friend", "colleague", "acquaintance", etc.), then the dyad can be said to be in a relationship of that type, governed by the extent of the overlap.

While a dyad's interaction history plays a strong role in establishing their relationship (the set of accommodation behaviors they will readily engage in), and provides content for continuity behaviors (referring back to shared experiences) and other functions, their history is not an intrinsic part of the representation of their relationship according to the above definition. To see why, note that dyads with relational expectations but no history (e.g. couples in an arranged marriage, business people in contractual business arrangements, etc.) are said to be in a relationship, at least from the first moment of acting in accordance to their relational expectations, whereas dyads with a history but who have no further expectations of relational interactions (e.g., deceased partner, real estate agent following a closing, divorcees, etc.) are said to no longer be in a relationship.

Finally, while the definition above states explicitly what two agents will do for each other, it also implicitly defines what they will *not* do with each other. If there is no expectation that G2 will accommodate G1 on α , then G1's proposing a collaboration on α is "marked" and must be carefully negotiated (discussed further in 3.7). This view allows common relational maxims such as "avoid intimate topics", "avoid impositions" or "avoid face threats" to be expressed by the exclusion of activities that would violate these rules from the set of activities sanctioned by the relationship.

3.5 Keeping Score

It is in an agent's best interest to know what it can rely on its relational partners for, so that it can plan its life accordingly. Unfortunately, except in certain formal situations (e.g., contractual relationships) this information has to be inferred because relationships are typically in a state of flux, and because rejections (finding out that a partner *won't* accommodate you on a specific activity) are significant blows to one's self esteem and represent "face threats" (Goffman, 1967).

An exact calculus of relational expectations has yet to be developed, but there are some general psycho-social principles that can be used. The most fundamental principle is that when a dyad engages in a particular collaborative activity, it increases each of their expectations that they can engage in the same activity again in the future, absent explicit evidence to the contrary. This implies that one way to change a relationship is to simply do something new together. This also implies that an existing relationship can be maintained (expectations kept high) simply through periodic performance of the collaborative activities that constitute the relationship.

Partners also perform explicit tests of the status of their relationship by proposing activities just to see if they get uptake or not (e.g., "why don't you come meet my parents?") or to simply confirm what they think the current status of the relationship is (Duck, 1991).

Relational stereotypes also play a significant role in relational expectations. To the extent that the set of accommodation behaviors defining a dyad's relationship overlaps with a socio-

cultural relational stereotype, their expectation of being able to perform any of the recipes indexed by the stereotype is increased. Thus, as a dyad begins to perform activities together, their expectations about continuing those activities *and similar activities* normally increase, where similarity here is partly governed by the stereotypes. For example, once a dyad begins to do many of the things that "friends" do (in a particular culture) they can expect to do all of the other things friends normally do together, absent reasons for thinking otherwise.

In addition to actual past history and relational stereotypes, expectations can be influenced by generalized notions of *trust* in one's relational partner.

3.5.1 Trust

Although agents may keep track of separate relational expectations for the entire range of accommodation behaviors they expect to collaborate on with their partner, it is useful (e.g., for planning purposes) to have generalized expectations about the likelihood of a partner meeting one's relational expectations. Such generalized expectations are called *trust* and, in social psychology, are generalized over all interactions with an individual ("specific interpersonal trust" (Johnson-George & Swap, 1982)), over all interpersonal relationships ("generalized trust", strongly related to attachment style (Berscheid & Reis, 1998)), or over all interactions in a particular context (e.g., professional vs. social relationships (Barber, 1983)).

Trust is a central concept in relationships, both because it is a prerequisite for any kind of collaborative activity, and because it is one of the primary outcomes of and reasons for engaging in relationship development. The literature on trust spans the disciplines of sociology, social psychology, and philosophy, and several computational models of trust have been developed in the software agents and distributed AI fields as well.

Relationally, trust is an antecedent to self-disclosure (Wheeless & Grotz, 1977) and is an important component of intimacy (Berscheid & Reis, 1998). Trustworthy sources are also generally more persuasive (at least for certain types of decisions) (Petty & Wegener, 1998).

Deutsch provides a detailed theory of trust, validated through a series of empirical studies of cooperation among human subjects (Deutsch, 1973). In his theory he suggests that trust involves an analysis of the benefits vs. the potential costs to the trustor for taking a particular action. He suggests that trust comes into play in those situations in which one makes a decision whose negative consequences potentially outweigh the positive consequences. He proposes a utility theoretic framework in which a trusting decision is made when the outcome of that decision is an ambiguous situation in which the following holds:

$$\begin{aligned}
 & Va^+ < Va^- \\
 & \text{and} \\
 & Va^+ \times S.P.^+ > Va^- \times S.P.^- + K
 \end{aligned}$$

Where,

Va^+ is an event having positive motivational significance (value, or benefit).

Va^- is an event having a negative motivational significance (value, or benefit).

$S.P.^+$ is the subjective probability of attaining Va^+ .

$S.P.^-$ is the subjective probability of Va^- .

K is the "security level" that the individual needs for action (trait).

The subjective probabilities involved in this assessment are based on many factors, but the most important ones are one's own past experiences in similar situations, the past experiences of others, the opinions held by others whom one respects, one's personal

assumptions about the benevolence-malevolence of the reality one is in, and one's confidence about being able to influence the occurrence of Va^+ or the nonoccurrence of Va^- through one's own actions or through available help. As an example, suppose I asked a close friend to mail an important letter for me. This is a trusting action on my part, because $Va^+ < Va^-$, where Va^- is the cost of the letter not getting mailed, but my assessment of $S.P.^-$ is sufficiently low with respect to my assessment of $S.P.^+$ that I decide to take this action anyway.

One of the most comprehensive computational models of trust was developed by Marsh to research models of cooperation among agents in Distributed Artificial Intelligence simulations (Marsh, 1994). Marsh investigated the role of trust in the initiation and maintenance of cooperation between agents. In his framework an agent x decides to cooperate with another agent y in situation α when:

$$T_x(y, \alpha) > \text{Cooperation_Threshold}_x(\alpha)$$

Where, $T_x(y, \alpha)$ is x 's trust in agent y for situation α , and is given by:

$$T_x(y, \alpha) = [U_x(\alpha) + \hat{T}_x(y)] \times I_x(\alpha) \times (\hat{T}_y(x))^x$$

And,

$U_x(\alpha)$ is the utility of α to x .

$\hat{T}_x(y)$ is x 's general trust in y .

$I_x(\alpha)$ is the importance of α to x .

$(\hat{T}_y(x))^x$ is x 's estimate of y 's trust in x .

The two estimates of trust are computed by integrating over past experience with the other agent (initialized with a "generalized" trust value). Marsh presents three simple ways of doing this--optimistic estimation (maximum), pessimistic estimation (minimum), and "pragmatic" estimation (averaged)--and also a method which takes into account a fixed memory of agents by averaging over a window of the last N interactions.

The Cooperation Threshold is defined as:

$$\text{Cooperation_Threshold}_x(\alpha) = \frac{\text{Perceived_Risk}_x(\alpha)}{\text{Perceived_Competence}_x(y, \alpha)} \times I_x(\alpha)$$

Where, Perceived_Risk is x 's perception of the risk associated with situation α , and $\text{Perceived_Competence}$ is x 's perception of y 's competence in performing in situation α .

While the exact form of Marsh's model is debatable, he has motivated his choice of variables and their direction of contribution in the estimation of trust and cooperation from extensive research into the literature on trust. It seems clear that trust plays a role in an agent's assessments of whether another agent will accommodate it or not, or whether it should accommodate another agent, and that the estimation of trust involves a cost/benefit analysis involving reasoning about past experiences with the other agent, the importance or utility of the outcome, and the risks involved. One drawback of using Marsh's work for a relational model is that it provides a generalized expectation; a calculus for determining the influence of

this on the specific expectations that an agent has for accommodation has yet to be determined.

3.6 Motivation to Meet Relational Expectations

Why should agents bother to act in accordance with the relational expectations they have established with their partner? One reason is to simply maintain the relationship so as to not lose any benefits that might later accrue from it. Many relationships, such as friendships, are voluntary, and violations of the “unwritten contract” between the partners can result in relationship dissolution (wiseman, 1986), resulting in loss of provisions to both partners.

3.6.1 Rights & Obligations

Another motivation for not violating relational expectations is the threat of retaliatory action by one’s partner and the society in general. These actions are sanctioned under the concepts of rights and obligations.

As Jackendoff observed, rights and obligations are fundamental to the fabric of human social organization and appear to be universal concepts in human societies (Jackendoff, 1999). Jackendoff’s analysis is directed at the question of how people conceptualize situations in which someone can be said to have a right or an obligation, and relies on the formal framework of “conceptual semantics”, which encodes regularities in people’s minds (rather than, e.g., formal statements about the real world). In this framework, rights and obligations are like facts-- “objectively determinable persistent entities”--that have their own logic which Jackendoff guesses might possibly be innate (aside from their universality, even primates seem to have some notion of a moral dimension of obligations). Rights and obligations can be either persistent and universal with respect to some class of actions, or existential in which, like intentions, once satisfied (an obligation fulfilled or a right exercised) they cease to exist. Rights have an inherently positive value, and are thus to be desired by an agent, whereas obligations have an inherently negative value.

A fundamental aspect of rights and obligations that distinguishes them from other kinds of modals (such as *may* and *should*) is the consequence of noncompliance. The inference rule for obligation noncompliance can be paraphrased as:

IF G2 has an obligation to do an action α for the benefit of G1 AND
G2 does not perform this action (in a suitable time-frame)
THEN
G1 has the (existential) right to do an action β that will have negative value for G2
in exchange for G2’s noncompliance

The parallel inference for rights is:

IF G1 has a right to perform an action α AND
G2 prevents G1 from performing α
THEN
G1 has the (existential) right to do an action β that will have negative value for G2
in exchange for G2’s violation

There are also moral dimensions to these phenomenon. In the case of unfulfilled obligations, all members of the society are morally justified in sanctioning someone for their noncompliance, whereas in the case of violated rights, everyone is morally justified in sanctioning the offender.

Traum also introduced an attitude of obligation in his work on modeling intentions in dialogue (Traum & Allen, 1994). In his model, an agent's behavior is determined by a number of factors, including the agent's current goals in the domain, and a set of obligations that are induced by a set of social conventions. The agent considers both goals and obligations, to the extent that it can, when planning actions. He used this framework to explain how and why an agent responds to dialogue moves by stating that certain moves obligate the listener (relative to some set of social norms) to respond. For example, when asked a question this creates an obligation in the listener to respond with an answer.

3.6.2 Implications

In light of Jackendoff's analysis, relational expectations of accommodation can be seen as persistent obligations to the agent who is to be accommodating, and persistent rights to the agent who expects to be accommodated. In this view, violating relational expectations opens up the possibility of retaliation by the thwarted partner and, especially to the extent that the relationship fits a socio-cultural stereotype, by other members of society. Relative to Traum's model, relational expectations actually fit more in line with his view of "social norms" in that they are persistent and give rise to specific collaborative behavior when the need arises.

3.7 Relational Dynamics

As defined above, a relationship not only defines what a dyad can do together, but implies that any activities not sanctioned by the relationship are forbidden. Given that most relationships are continuously evolving, how do changes in relational expectations occur? They must be carefully negotiated. Explicit forms of negotiation can be, in themselves, very subtle and complex. However, relational negotiation is even more 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.

3.7.1 Negotiation

The basic form of negotiation occurs when one agent makes a proposal for a new activity and the other agent accepts or rejects it. Human relational negotiations are typically very subtle and much more complex. However, even when negotiation moves are entirely explicit, negotiation is still a very complex phenomenon. Sidner defines an artificial language within which negotiations can take place, and includes actions such as AcknowledgeReceipt (of a proposal), Counter (a counter-proposal), and RetractProposal in addition to the basic propose and accept actions (Sidner, 1994).

3.7.2 Face

One theory that sheds a lot of light on the forms in which relational negotiation must occur is Goffman's notion of "face". This construct came out of work in sociology (Goffman, 1967) and has been most productively employed in the field of sociolinguistics (Brown & Levinson, 1987). It has been shown to govern a significant amount of behavior in social

interactions, and is itself governed by the nature of the relationship between the interactants. "Facework" (strategies to manipulate this construct) can be used to change the nature of a relationship, and is thus relevant to the design of relational agents.

In Goffman's dramaturgical approach to social interaction, he defined an interactant's "line" as the patterns of action by which an individual in an interaction presents an image of himself and the situation, that is his social role in the current joint activity (Goffman, 1967). The notion of "face", Goffman went on to say, is "the positive social value a person effectively claims for himself by the line others assume he has taken during a particular contact". An interactant maintains face by having their line accepted and acknowledged by their interactants. Events which are incompatible with their line are "face threats" and are mitigated by various corrective measures if they are not to lose face.

Brown and Levinson used Goffman's notion of face in the formulation of their theory of the use of politeness forms in language (Brown & Levinson, 1987). They defined positive face as an individual's desire to be held in esteem by their interactants, and negative face as an individual's desire for autonomy, and went on to characterize the degree of face threat of a given speech act as a function of power, social distance, and the intrinsic threat (imposition) imposed by the speech act. Brown and Levinson's theory accounts for some cross-cultural universals in sociolinguistic behavior by explaining when individuals are likely to use different forms of politeness in making a request, given the nature of the request and the relationship between the interactants.

Relative to my theory of relationship, facework can be seen as predicting the kind of mitigating action required when a change in relationship is attempted. When actions sanctioned by the relationship are performed, no mitigating action is required (and, via accommodation, no negotiation of any kind need occur). However, when actions outside the relationship are proposed, different kinds of mitigation are called for, from conventional politeness to going "off record" with a request and thereby making it very easy for the helper to reject the proposal without threatening the requestor's self-esteem. The degree to which the action is unexpected, relative to the existing relationship, determines the form of mitigation required (per Brown & Levinson) and the dimensions along which this degree of unexpectedness are measured are, according to Brown & Levinson, power and solidarity, the two primary dimensions used in dimensional models of relationship.

Facework normally plays out in the micro-structure of face-to-face conversation. Examples of strategies that can be employed include: hedged or indirect requests ("You wouldn't possibly want to go to the movies, would you?"); pre-requests ("Do you like movies?"); pre-invitations ("What are you doing this evening?"); and pre-announcements ("You know what I'd like to do?") (Levinson, 1983). Rejections are almost always indirect and often nonverbal, including such behaviors as pausing (allowing the proposer to retract their suggestion), gazing away, preface markers ("Uh", "Well"), and affective facial displays (rejections are a type of "dispreferred" conversational action (Levinson, 1983)).

3.7.3 Other Forms of Change

Another method of changing relational expectations is for one partner to simply start acting as if the change had already occurred. For example, Lim posits that when an individual is attempting to increase closeness in a relationship they will perform less than expected facework for "normal and familiar" joint activities, while those who are trying to increase

distance in a relationship will use facework strategies normally used in more distant relationships (Lim).

When one partner starts acting as if a change had occurred and the other goes along with it, the change behavior itself is a form of accommodation, and may lead both partners to believe that this method of change can be used again in the future.

3.8 The Affective Dimension

Much has been said about social dimensions of relationship, but not much about the emotional aspects thus far. However, emotions play a crucial role in all human relationships. Relative to the relational theory presented above, emotions (e.g., love, happiness), their grounding (through empathy), and management (through emotional support) are fundamental provisions in most close personal relationships. Further, deviations from relational expectations give rise to some of the strongest emotions humans feel (e.g., shame, guilt, embarrassment, jealousy, and social anxiety) and thus serve a regulatory role in helping to ensure that partners meet each other's expectations.

3.8.1 Affective Relational Provisions

Emotions (e.g., love, joy) and behaviors that help manage emotions (e.g., empathy) are among the most important provisions of close personal relationships. Many kinds of relationship stereotypes, such as passionate love, are defined in terms of the emotions felt by one partner towards the other (Brehm, 1992).

According to Goleman, being able to manage emotions in someone else is the core of the art of handling relationships (Goleman, 1995), and the ability to provide emotional support is also frequently mentioned as one of the requisite characteristics of a good friend or intimate partner (Cole & Bradac, 1996; Gill, Christensen, & Fincham, 1999; Goldsmith, McDermott, & Alexander, 2000; McGuire, 1994). Emotional communication between partners is a crucial ingredient in most relationships. Relationship quality is often defined in terms of the quality of emotional communication between partners; "it is taken for granted that marital harmony is strongly related to effective communication between spouses, and that effective communication, to some optimal degree, involves spouses' understanding of each others' thoughts and feelings" (Berscheid & Reis, 1998). Empathy--the process of attending to, understanding, and responding to another person's expressions of emotion--is a pre-requisite for providing emotional support which, in turn, provides "the foundation for relationship-enhancing behaviors, including accommodation, social support, intimacy, and effective communication and problem solving" (Berscheid & Reis, 1998) (Okun, 1997).

In addition to being important in intimate relationships, emotional support provisions are also crucially important in most helping professions. The most significant empirical support of this phenomenon is in the field of psychotherapy, in which measures of "working alliance"--the trust and belief that the therapist and patient have in each other as team-members in achieving a desired outcome--show consistently high correlations with successful outcomes (Horvath & Symonds, 1991). Even in physician-patient interactions, physician empathy for a patient plays a significant role in prescription compliance, and a physician's *lack* of empathy for a patient is the single most frequent source of complaints (Frankel, 1995).

3.8.2 The Regulatory Role of Affect

Relationships play a central role in the emotional life of all humans. According to Lazarus "most emotions involve two people who are experiencing either a transient or stable interpersonal relationship of significance" (Lazarus, 1994), and Bowlby contends that most intense emotions arise when people are forming, maintaining, disrupting, terminating, or renewing close relational ties with others (Bowlby, 1979). Many emotions, such as jealousy or passionate love, can only occur within a relationship, while other emotions, such as loneliness, occur because of problems with or lack of a desired social relationship.

One function of these strong emotions is to serve as basic regulatory mechanisms to ensure that relational expectations are met. Violations of relational expectations give rise to emotions such as shame, guilt, embarrassment, and social anxiety and act as a motivator for the partner feeling them to take mitigating action.

3.8.3 Affect in Relational Negotiation

Emotions also play an important role in the negotiation process of relationship building and maintenance. As stated above, relational negotiation is usually conducted in a tacit, off-record manner, with the proposals and rejections handled in an indirect manner to save face. Uptakes and rejections, in particular, are often achieved through positively and negatively valenced emotional displays (e.g., happiness display for uptake, disgust, contempt or anger displays for rejection).

3.9 Application to the Micro-Structure of Face-to-Face Conversation

Face-to-face conversation is the primary modality used to build and maintain human relationships. It should be apparent from the definitions above that negotiation of relational expectations, for most types of relational provisions, would be very difficult to accomplish without language. According to Duck, "Language—or more broadly, communication and all that is culturally encoded within it—is ... a crucial basis for establishing ways in which we conduct human relationships and judge their quality" (Duck, 1998). In other work, Duck even defines relationships in terms of "shared meaning systems" whereby the individuals in the relationship come to see and evaluate events in the same way, principally through the use of language (Duck, 1995).

Language is also required in order to enact many kinds of relational provisions, such as informational support, and it seems to be especially important in emotional support processes. For example, there is much evidence that talking about traumatic events helps people to deal with them. According to Burleson and Goldsmith, all of the emotions, images and thoughts associated with a trauma are held together by the traumatic incident itself, and talking about the event may help dismantle the phenomenal whole that traumas constitute (Burleson & Goldsmith, 1998). The very process of putting the event into words, subject to the syntactic and pragmatic constraints of language and the need to decontextualize it to describe it to another, helps the distressed person get a new perspective on the event and reappraise it. Another series of studies showed that when people fail to articulate a traumatic event verbally, they fail to deal with it effectively. Thus, conversation seems to play a central role in coping with emotion. (Andersen & Guerrero, 1998b).

Not all forms of language use, however, are equally effective for relational communication. Face to face conversation is the primary, universal mode of human-human

communication. According to Fillmore, "the language of face-to-face conversation is the basic and primary use of language, all others being best described in terms of their manner of deviation from that base" (Fillmore, 1975), and Levinson, "face-to-face interaction is not only the context for language acquisition, but the only significant kind of language use in many of the world's communities, and indeed until relatively recently in all of them" (Levinson, 1983). Face-to-face conversation is particularly important for relational communication because of the availability of nonverbal cues, such as facial expressions, body posture and hand gesture.

3.9.1 The Role of Nonverbal Behavior

According to Argyle, nonverbal behavior is used to express emotions, to communicate interpersonal attitudes, to accompany and support speech, for self presentation, and to engage in rituals such as greetings (Argyle, 1988). Next to coverbal behavior and emotional displays, the most important use of nonverbal behavior in relational dialogue is the display of interpersonal attitude. The display of positive or negative attitude can greatly influence whether we approach someone or not and our initial perceptions of them if we do.

The most consistent finding in this area is that the use of nonverbal "immediacy behaviors"--close conversational distance, direct body and facial orientation, forward lean, increased and direct gaze, smiling, pleasant facial expressions and facial animation in general, nodding, frequent gesturing and postural openness--projects liking for the other and engagement in the interaction, and is correlated with increased solidarity (perception of "like-mindedness") (Argyle, 1988; Richmond & McCroskey, 1995). Other nonverbal aspects of "warmth" include kinesic behaviors such as head tilts, bodily relaxation, lack of random movement, open body positions, and postural mirroring and vocalic behaviors such as more variation in pitch, amplitude, duration and tempo, reinforcing interjections such as "uh-huh" and "mm-hmmm", greater fluency, warmth, pleasantness, expressiveness, and clarity and smoother turn-taking (Andersen & Guerrero, 1998a).

The importance of nonverbal behavior for relational communication is also supported by studies comparing the affordances of different communication media. Several studies have found that the additional nonverbal cues provided by video-mediated communication do not affect performance in task-oriented interactions, but in interactions of a more relational nature, such as getting acquainted, video is superior (Whittaker & O'Conaill, 1997). These studies have found that for social tasks, interactions were more personalized, less argumentative and more polite when conducted via video-mediated communication, that participants believed video-mediated (and face-to-face) communication was superior, and that groups conversing using video-mediated communication tended to like each other more, compared to audio-only interactions. The importance of nonverbal behavior is also supported by the intuition of businesspeople who still conduct most important business meetings face-to-face rather than on the phone.

3.9.2 Relational Communication

Relative to the theory outlined above, conversation can be seen as an elaborate dance in which every utterance made by one partner contributes to his assessments of relational expectations and trust in the other, and thus to his perception of the relationship. According to Duck, every verbal message contains two elements: 1) propositional content and 2) a message

about the relationship; “You can barely utter a word without indicating how you feel about the other” (Duck, 1998).

Every utterance must either fit squarely within the relational expectations that partners have (serving to ratify the relationship by keeping current expectations high), or is to some extent unexpected, in which case it should be marked and/or mitigated to avoid threatening the face of one or both partners. Politeness theory provides some guidelines about the form these marking and mitigating strategies should take (Brown & Levinson, 1987).

This is not to say that partners should use the exact same language in every situation, but that at some level of abstraction the kinds of topics they discuss and the activities they engage in with language has a regularity to it that is defined by their relationship. The exact form of language used is, in fact, expected to vary subtly with the context (both immediate and historical), and conversationalists who do not adapt their language in this way are seen as inept and un-attentive to the conversation (Duck, 1991).

3.9.3 Social Deixis

One way in which language can be used to set relational expectations is through social deixis, or what Svennevig calls “relational contextualization cues” (Svennevig, 1999), which are “those aspects of language structure that encode the social identities of participants...or the social relationship between them, or between one of them and persons and entities referred to” (Levinson, 1983). Politeness strategies fall under this general category (facework strategies are partly a function of relationship), but there are many other language phenomena which also fit, including honorifics and forms of address. Various types of relationship can be grammaticalized differently in different languages, including whether the relationship is between the speaker and hearer as referent, between the speaker and hearer when referring to another person or entity, between the speaker and bystanders, or based on type of kinship relation, clan membership, or relative rank (Levinson, 1983). One of the most cited examples of this is the *tu/vous* distinction in French and other languages.

For English, Laver encoded the rules for forms of address and greeting and parting as a (partial) function of the social relationship between the interlocutors, with titles ranging from professional forms (“Dr. Smith”) to first names (“Joe”) and greetings ranging from a simple “Hello” to the more formal “Good Morning”, etc (Laver, 1981).

Forms of language may not only reflect existent relational status, but may be used to negotiate changes in the relationship, by simply using language forms that are congruent with the desired relationship. As stated above, Lim observed that partners may change their facework strategies in order to effect changes in the relationship (Lim). And, according to Svennevig:

The language forms used are seen as reflecting a certain type of relationship between the interlocutors. Cues may be used strategically so that they do not merely reflect, but actively define or redefine the relationship. The positive politeness strategies may thus ... contribute to strengthening or developing the solidarity, familiarity and affective bonds between the interactants. The focus is here shifted from maintaining the relational equilibrium toward setting and changing the values on the distance parameter (Svennevig, 1999, pg. 46-47).

3.9.4 Social Dialogue

Common examples of social dialogue are small talk, conversational storytelling, gossip, getting acquainted talk, and joke-telling. What do these speech genres have in common? It is easiest at this point to define social dialogue by what it is not; it is talk whose focus is not on factuality and instrumentality, in short, talk that is not task-oriented. Given that all utterances carry relational meaning (as described above), what social dialogue is focused on is primarily the negotiation of the interpersonal relationship between the interlocutors.

Perhaps the purest form of social dialogue is what Malinowski referred to as "phatic communion", "a type of speech in which ties of union are created by a mere exchange of words." This is the language used in free, aimless social intercourse, which occurs when people are relaxing or when they are accompanying "some manual work by gossip quite unconnected with what they are doing" (Malinowski, 1923). An example of this type of small talk is the situation in which two strangers approach each other and the speaker (S) says to the hearer (H):

(1) Lovely day.

Utterances of this type have not received adequate treatment in linguistics or the philosophy of language. In these fields, "meanings" are taken to be those properties of utterances or texts which have their impact on hearers, such as describing states of the world or predicting or explaining human behavior (Devitt & Sterelny, 1999). The predominant, classical approach to meaning is that the meaning of a declarative statement is its truth-value when compared against the world. However, this captures very little of the meaning of utterance (1) in the situation described. The impact of this statement would be roughly the same if S made it under a very wide range of weather conditions, most of them far from ideal (although, if the weather is plainly very poor it could have an added element of irony associated with it). Thus, the classical notion of truth-value seems to play very little role in the meaning conveyed. In many situations, the impact of this statement on H would also be roughly the same if S approached H with a warm smile and gaze and uttered 'Lovely X', where 'X' could be composed of almost any short sequence of phonemes pronounceable by S. Thus, the meaning conveyed by this utterance in this kind of situation also has very little to do with its propositional content in the first place. The classical approaches to meaning do not help much in explaining the meaning of utterances in social dialogue.

As Wittgenstein and others have pointed out, there are many kinds of meaning. In addition to the classical, propositional notion of meaning described above, I claim there are two other kinds of meaning that are important for conversation: interactional and relational. Interactional meaning is that aspect of an utterance which affects what interlocutors do in the current interaction, affecting their behavior in such things as turn-taking, grounding, and emphasis (Duncan, 1974). Relational meaning is that aspect of an utterance which indexes the relationship between the interlocutors, and the rights and obligations entailed. The 'truth' of this meaning is the degree to which these obligations are adhered to (the degree of authenticity or sincerity).

Returning again to utterance (1), it can be seen to have: a propositional meaning (some congruence with the current weather conditions); an interactional meaning (transitions into a

conversation (Laver, 1975)); and a relational meaning (indexed a particular type of relationship) all at once. Thus, "small talk" can be seen as a conversational recipe for a particular kind of talk, and is thus fundamentally an interactional construct. One of its primary functions is to provide the opportunity for social dialogue, since it is a type of talk in which there is minimal commitment to factuality and minimal obligations entailed. However, its use does index a particular type of relationship; stereotypically one between strangers or acquaintances who are on polite (if not "friendly") terms with each other.

Researchers in linguistics, sociolinguistics and social psychology have hypothesized various functions that social dialogue might have. As mentioned above, Malinowski observed that some forms of social dialogue (i.e., "phatic communion") are used to "establish ties of union" or "to establish bonds of personal union between people brought together by the mere need of companionship" (Malinowski, 1923). Jacobson included a "phatic function" in his well-known conduit model of communication, which is focused on the regulation of the conduit itself (as opposed to the message, sender, or receiver) (Jacobson, 1960). Laver concluded that the underlying goals of phatic talk are to establish relationships and achieve transition (Laver, 1975). Coupland & Coupland hypothesized that phatic discourse involves a minimized commitment to factuality, open disclosure and seriousness, and that utterances could be classified along a continuum of "phaticity", rather than being discretely phatic or not (Coupland, Coupland, & Robinson, 1992). Dunbar found that the purpose of social dialogue is primarily to build rapport and trust among the interlocutors, provide time for them to "size each other up", establish an interactional style, and to allow them to establish their reputations (Dunbar, 1996).

In summary, social dialogue is conversation in which issues related to any instrumental task talk is backgrounded and interpersonal issues are (by default) foregrounded. Relative to the theory presented above, then, the role of social dialogue is primarily to build or maintain relational expectations and thereby trust in the relational partner.

3.9.5 A More General Model of Facework in Face-to-Face Conversation

In this section I present a dimensional model of relationship and a new model of facework that is used as the basis for a computational model of dialogue planning (presented in Chapter 4) and an evaluation of the effects of social dialogue on trust in relational agents (presented in Chapter 5). This work was previously published in (Cassell & Bickmore, 2001).

As discussed in section 3.1, there is an equivalence between dimensional and provisional models of relationship, in that dimensional models provide an abstract feature space within which specific provisional models can be placed. The dimensional model used in this theory of facework, while not providing the granularity of specific beliefs and intentions of the interlocutors with respect to their relationship, provides a good starting place for explorations in generating dialogue moves that are not only congruent with an existing relationship but serve to move the relationship in a desired direction.

Prior Models of Facework

As described in section 3.7.2, Brown and Levinson extended Goffman's notion of face in their theory of politeness forms in language (Brown & Levinson, 1978). They defined positive face as an individual's desire to be held in esteem by his interactants, and negative face as his desire for autonomy, and characterized the degree of face threat of a given speech

$$\text{face threat} = f(\text{SA}_{\text{intrinsic}}, \text{Power}, \text{Distance})$$

$$\text{SA}_{\text{intrinsic}} = f(\text{SA})$$

Where,

SA_{intrinsic} = the intrinsic threat of the speech act
 SA = denotes a class of speech acts
 Power = power relationship between speaker and hearer

Figure 3-1. Brown & Levinson's Face Threat

act as a function of power, social distance, and the intrinsic threat (imposition) imposed by the speech act. That is, the face threat to the hearer can be given by the formula in Figure 3-1.

The 'intrinsic threat' parameter accounts for the fact that certain speech acts are more of a threat than others. For example, an informing is less of a threat than a request for information which is less of a threat than a rejection. Distance is defined to be "a symmetric social dimension of similarity/difference within which the speaker and hearer stand for the purposes of this act" and power is the ability of one interactant to control the behavior of the other.

If a significant threat will result from the speaker producing the indicated speech act, then the speaker has several options: 1) don't do the act; 2) do the act "off record"; 3) do the act "on record" with redressive action (negative politeness strategies); 4) do the act on record with redress action (positive politeness strategies); 5) do the act on record, "baldly". Following Grice's (Grice, 1989) description of how to fail to fulfill the conversational maxims, these options are ranked in order of decreasing ability to mitigate a threat, thus the most threatening acts shouldn't be done at all, while the least threatening acts can be done baldly on record. Examples of "off record" acts are hinting and/or ensuring that the interpretation of the utterance is ambiguous (e.g., "I'm thirsty."). Negative politeness strategies include those which are oriented towards the autonomy concerns of the listener (e.g., "Could you bring me a drink?"), while positive politeness strategies address the esteem concerns of the listener (e.g., "Hey my friend, get me a drink.>").

Svennevig extended Brown and Levinson's model by noticing that the threat perceived from different types of speech acts can change based on context, and in particular based on the relationship between the speaker and hearer (Svennevig, 1999). For example, close friends have established a set of mutual rights and obligations and thus do not experience certain acts (such as requests) as face threatening, but rather as confirming and reestablishing their relational bonds. (This view actually fits the theory presented in this chapter very well.) Thus, his extension to the model can be characterized as shown in Figure 3-2.

$$\text{face threat} = f(\text{SA}, \text{Power}, \text{Solidarity}, \text{Familiarity}, \text{Affect})$$

Where,

SA = denotes a class of speech acts (not mapped to an "intrinsic" threat value)

Figure 3-2. Svennevig's Face Threat

Svennevig's dimensional model of interpersonal relationships has four dimensions:

- Power - the ability of one interactant to control the behavior of the other.
- Solidarity – degree of "like-mindedness" or having similar behavior dispositions (e.g., similar political membership, family, religions, profession, gender, etc.), and is very similar to the notion of social distance used by Brown and Levinson (Brown & Levinson, 1978).
- Familiarity - based on social penetration theory (Altman & Taylor, 1973), which claims to account for the establishment and growth of interpersonal relationships, this dimension describes the way in which relationships develop through the reciprocal exchange of information, beginning with relatively non-intimate topics and gradually progressing to more personal and private topics. The growth of a relationship can be represented in both the breadth (number of topics) and depth (public to private) of information disclosed.
- Affect - the degree of liking the interactants have for each other, and there is evidence that this is an independent relational attribute from the above three (Brown & Gilman, 1989).

A New Model

Based on analyses of interactions involving both task and social dialogue (initial meetings between real estate agents and buyers), I have extended Svennevig's model of face threat to include measures of topic coherence and topic intimacy.

It is clear that the introduction of conversational topics that are at a significantly deeper level of familiarity than is expected relative to the existent relationship and activity will be seen as a face threat. For example, if a stranger on the street asked you how much money you had in your bank account, you would likely perceive this as a threat to your face. Such a kind of face threat is key to task encounters where strangers must interact, and occasionally share personal information. I term this a "Social Penetration" threat, or SP_{threat} .

Topics that are at the expected level of familiarity but which are completely unrelated to the topic at hand also seem to be face threats, but have not been accounted for in a general way in previous theory. While a subset of these have been addressed in Brown and Levinson's theory (e.g., rejections), moves which are deemed dispreferred based solely on their sequential placement in conversation cannot be accounted for, given Brown & Levinson's use of isolated speech acts as their point of departure. Instances of such "sequential placement" threats are failing to demonstrate the relevance of a conversational story, failure to appreciate conversational stories following their conclusion (Jefferson, 1978), or introducing conversational topics or stories which are not related to the on-going discourse (not "locally occasioned" (Sacks, 1995)). Thus, for example, if you tell someone a long humorous story, and they respond, not by nodding or otherwise acknowledging your story, but instead by changing the topic, that will threaten your face.

The resulting model of face threat is shown in Figure 3-3. It is novel in that it goes beyond the analysis of a single speech act to acknowledge and incorporate the role of discourse context into the determination of face threat. This model can be used by a relational agent to provide information about dialogue moves that it should or should not make in order to avoid threatening the face of its users.

face threat	= $f(\mathbf{SA}_{\text{threat}}, \mathbf{SP}_{\text{threat}})$
SA_{threat}	= $f(\mathbf{SA}_k, \{\mathbf{SA}_1, \dots, \mathbf{SA}_j\}, \text{Power, Solidarity, Familiarity, Affect})$
SP_{threat}	= $f(\text{FamiliarityDepth, TopicDepth})$
Where,	
SA _{threat}	= Threat due to the speech act.
SP _{threat}	= Threat due to violation of social penetration theory.
SA _k	= The class of speech act.
{SA ₁ , ..., SA _j }	= The discourse context of speech acts into which SA _k will be introduced. For example, SA ₁ could represent the overall conversation, and SA _j represents the activity which SA _k will become a constituent of.
TopicDepth	= The "depth" of the topic to be introduced (wrt social penetration theory).
FamiliarityDepth	= The current value of the depth dimension of the familiarity aspect of the relationship among the interlocutors.

Figure 3-3. A New Model of Face Threat

3.9.6 Face-to-Face Conversation with Relational Agents

It should be clear that emulations of face-to-face conversation are the ideal medium for relational agents to use in order to build complex relationships with their users. The model of face threat above can be used on an utterance-by-utterance basis by a relational agent to inform its selection of next topics in this setting. Of course, many other factors need to be taken into account in selecting the content and form of a next utterance, including the attentional and intentional state of the dialogue (Grosz & Sidner, 1986), information structure (Prince, 1981), and the goals of the agent. Taken together, these factors allow an agent to meet the relational expectations of a user and thus build trust simply by engaging in dialogue with them. Chapter 4 presents a computational model that applies this model of facework to the problem of planning task and social dialogue in a particular type of service encounter, and Chapter 5 presents the results of an evaluation of the effectiveness of this model.

3.10 Application to the Macro-Structure of Long-Term Relationships

Relational partners cannot typically spend all of their time together. Given that actual enactment of relational provisions is the primary means of maintaining relational expectations, these expectations may decay over time when partners are apart, may be challenged by information from third parties, or may be subject to doubting brought about through cognitive re-appraisal processes. However, even with these potential threats to relationship longevity, most relationships seem to survive times when partners are apart, whether it be hours, days or even years. Something persists in their heads which lets them pick up where they left off the last time they were together. According to Duck:

We do not go through a ritual each breakfast time where we treat each other like strangers and run through a whole range of rewarding techniques to re-establish the relationship and take it to where it was the day before: we behave that mental way only with friends we have not seen for ages. The remarkable fact about daily life is that continuities exist in our minds and do not have to be worked for, once the relationship is defined and established (Duck, 1998).

How do partners maintain their relationships over the long term and especially when they are apart? What exactly do they need to remember about their prior interactions when they meet again? And, how do relationships normally change over time? Relational agents that are designed to interact with users for more than a single brief session—and especially those which live with their users for months or years—need to be designed with these issues in mind.

3.10.1 Relationship Maintenance

A good question phrased by Duck is, “Do relationships fall apart unless they are maintained, or do they stay together unless they are taken apart?” (Duck, 1988). While some relationship theorists weigh in on both sides of this question, it is clear that there are specific behaviors that partners engage in which serve to actively maintain their relationship.

Partners perform a wide range of behaviors in order to keep their relational expectations high. Following the theory presented above, the most fundamental way that partners do this is by simply performing the collaborative actions which constitute their relationship. This is supported by research finding that the most common maintenance behavior performed by couples is simply sharing tasks (Dainton & Stafford, 1993).

In addition to routine maintenance behaviors, many researchers have investigated “strategic maintenance behaviors,” which are those performed with the conscious intent of maintaining the relationship. Some of the most common categories of these behaviors include prosocial behaviors (e.g., telling a joke, initiating interaction), ceremonial and ritual behaviors, meta-relational communication (talking about the relationship), and assurances (explicit statements about commitment to the relationship) (Gilbertson, Dindia, & Allen, 1998) (Dainton & Stafford, 1993; Dindia, 1994; Stafford & Canary, 1991)

Routine maintenance behaviors, on the other hand, are those behaviors performed on a regular basis without the conscious intent of maintaining the relationship, but which serve this purpose nevertheless. Gilbertson, et al, investigated the use of “continuity behaviors” among cohabitating couples, which are behaviors enacted before, during and after an absence to bridge the gap in the relationship caused by the absence (Gilbertson et al., 1998). Prospective behaviors (enacted before an absence) includes such things as telling one’s partner what you will be doing during the time apart, affirmations, and farewells; introspective behaviors (enacted during an absence) includes communicating with one’s partner (e.g., via telephone), and talking about them to others; and retrospective behaviors (enacted after an absence) included greeting and asking what your partner did during the absence.

3.10.2 Expected Trajectories of Relational Change

In addition to simply maintaining relationships, people generally expect relationships to change over time, in predictable patterns. According to Duck (Duck, 1988) as a relationship develops between individuals, the following things change:

- The content of interactions and the activities the partners engage in change.
- The diversity of interactions (the number of different things partners do together) generally increases.
- The qualities of interactions change, such as language forms and nonverbal behavior in face-to-face interaction.
- The relative frequency of interactions increases.
- Reciprocity decreases and complementarity increases. Reciprocity is when partners respond to each other's actions in similar ways (e.g., feeling obliged to repay social goods in kind); on the other hand, complementarity is when partners respond in different ways that fit together (e.g., one dominates and the other submits).
- Intimacy (physical and psychological) increases.
- Interpersonal perception aligns with self-perception. Partners' perceptions of themselves tend to align with their partner's perceptions of them.
- Commitment increases. Commitment is intention to continue in a relationship, and is a general measure of the strength of a relationship.

Several researchers have proposed "stage models" of relationship, which assume there are a fixed set of stages that different types of relationships go through. Examples of stage models are those by Lewis (Lewis, 1972), who proposes that (marital) relationships begin with similarity-based attraction then proceed through stages of good rapport, mutual self-disclosure, empathic understanding of the other person, role compatibility, and finally commitment to the relationship and identity as a couple. In another model, Reiss proposes that relationships go through four stages: initial rapport; mutual self-revelation; mutual dependency; and personal need fulfillment (Reiss, 1960). Stage models are now generally considered to provide very weak predictive power given their assumption of a fixed sequence of stages, since actual relationships often jump around among various stages in a non-linear manner (Brehm, 1992).

3.10.3 The Role of Personal History

According to the theory above, knowledge of past times together is only important for informing assessments of expectations of future accommodation behaviors; once these assessments are complete this knowledge can be discarded. Yet memory of specific past interactions seems to be very important to relational partners. For example, Planalp found that mutual knowledge and talking about past events were among the primary features that observers used to reliably differentiate between conversations involving friends and those involving acquaintances (Planalp, 1993). Exactly what role do these memories play, and do they represent an important feature to build in to relational agents who will be interacting with their users over long time spans?

As mentioned above, relational expectations can decay, especially in the absence of one's partner. Thus, one important role of remembering, and talking about, past activities together, is to keep relational expectations high. For example Stafford and Canary found that one of the maintenance strategies couples use is to remind each other about relationship decisions made in the past (Stafford & Canary, 1991).

As relationships develop, it is clear that partners gain knowledge about each other, both through informing and sharing experiences; increasing “common ground”.¹

This is not to say that partners’ memories of past shared experiences are perfect, or that they even agree on them. Even events as important as first meetings and the ways in which a relationship started (“origin stories”) are often remembered very differently by partners early in the relationship (Duck, 1991). Miell, in a study of weekly assessments of relationship strength over several months, found evidence to support the view that “memory for relationships ... as being an active, interpretive process, where the events of the past are reconstructed, rather than reproduced, in the context of current events” (Miell, 1987).

In conclusion, memory of one’s relational partner and specific past activities together is an important part of relationships, in that this memory provides a “common ground” or “shared meaning system” which lets dyads communicate more efficiently, re-evaluate their relational expectations, and provides fodder for certain maintenance activities such as reminding each other of past shared activities or relational decisions. Retention of this knowledge could be viewed as a type of accommodation behavior; partners routinely expect each other to interpret their utterances in part by filling in what is presumed to be shared knowledge. However, I still contend that relational memory is not an intrinsic part of the concept of relationship; you can know a lot about someone without having a relationship with them and you can have a relationship with someone (e.g., contractual) without any relational memory of past interactions.

3.10.4 Persistent Relational Agents

In order to build relational agents capable of multiple interactions with users over long periods of time it is clear that they should have at least some of the capabilities outlined above, including: some memory of information about the user and past interactions; some changes in relational behavior over time in ways expected by the user; and the use of routine and strategic relationship maintenance behaviors. Chapters 7-9 present the design and evaluation of a relational agent with these capabilities.

3.11 Summary

In this chapter I have outlined a theory of what a relationship is, some of the strategies that relational agents should use in making relational decisions, and a range of relational behaviors that can be used by them to build, maintain and evolve relationships with their users. I also discussed how this theory could be applied to the design and evaluation of relational agents in two realms: the micro structure of face-to-face conversation and the macro structure of long-term relationships. In the following chapters I will present implementations and evaluations of relational agents in each of these areas.

¹ Clark defines the “common ground” of two interlocutors as the sum of their mutual knowledge, mutual beliefs, and mutual suppositions (Clark, 1992). My use of this term is slightly different. Here, I am interested in the mutual knowledge that relational partners persistently carry over time, particularly *between* interactions.

Chapter 4

DISCOURSE PLANNING FOR RELATIONAL DIALOGUE

Social dialogue--crucial for building and maintaining relationships--is a very messy business when viewed from the logical framework of state-of-the-art discourse planning and text generation work in computational linguistics. Issues such as solidarity, common ground, liking, trust and politeness are at least as important as more traditional concerns such as discourse segmentation, presupposition and coherence. This chapter presents a computational model of mixed task and social dialogue that addresses some of these relational issues. The work presented here is reported in (Bickmore & Cassell, 2001) and (Cassell & Bickmore, 2001).

4.1 Introduction

As discussed in the previous chapter, most human relationships are built and maintained using language; thus, a relational agent must be able to have mastery over the relational aspects of language use in order to be effective. While it is possible to construct relational agents that rely solely on scripted dialogue and invariant surface forms for utterances, ultimately relational language must be planned and generated to achieve maximum generality, adaptability and variability in agents that are to serve in non-trivial roles, especially over long periods of time.

Planning relational dialogue is a very challenging problem. Conversation presents extreme time constraints on production which must be adhered to, lest the delays themselves carry unintended meaning (e.g., slow speech rates and inter-turn delays carry connotations of low intelligence in humans). Utterances must, at a minimum, be congruent with the assumed current state of the user-agent relationship, otherwise they may actually damage the relationship (e.g., by threatening the face of the user by using incorrect forms of politeness). Planning utterances that can actually move the relationship along a desired trajectory represents yet another level of complexity. Finally, assessing the ‘score’—the current status of the user-agent relationship—is itself an extremely difficult problem, ultimately requiring integration of verbal and nonverbal cues, but one which must be solved to provide the relational basis that is used for dialogue planning.

4.1.1 Problem: Dialogue Move Sequencing to Achieve Task Goals under Relational Constraints

This chapter presents a dialogue planner that solves a simplified version of the relational dialogue planning problem, while still addressing many of the important issues. The planner assumes an agent-initiated dialogue (a reasonable assumption for the problem domain described below), and dynamically sequences dialogue moves on the basis of discourse and relational context to satisfy the task goals of the agent given a set of relational constraints. The planner has at its disposal a range of both task moves (e.g., questions and statements about the task) and relational moves (primarily small talk, off-task remarks, questions and stories). Full

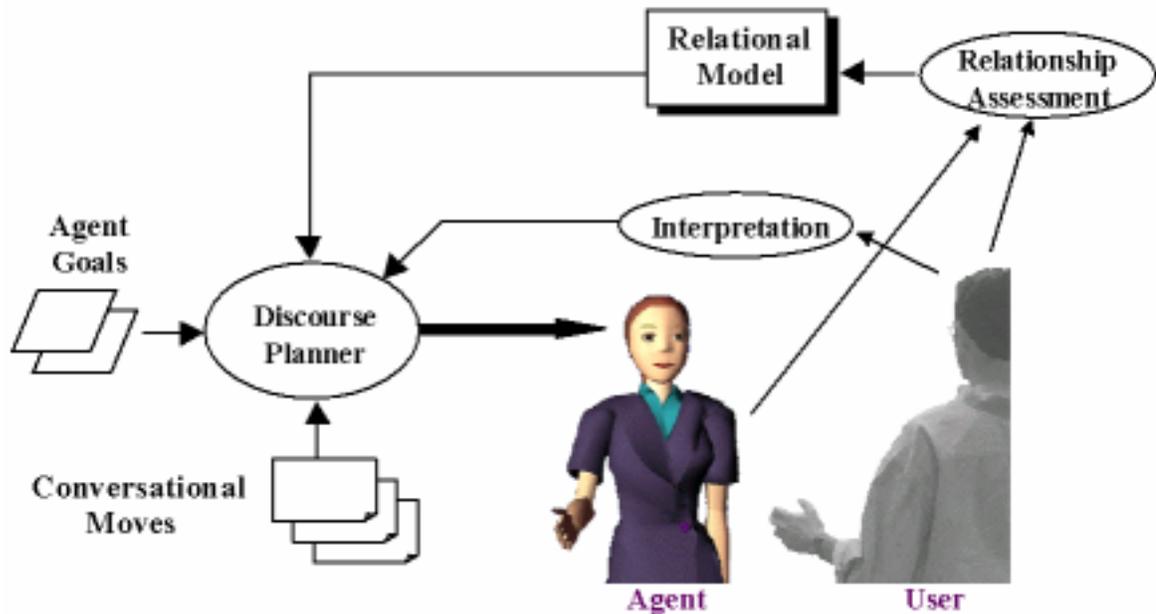


Figure 4-1. Relational Agent Discourse Planner Framework

text generation is not performed, but some surface variation is dynamically performed, such as the insertion of discourse markers at topic shift boundaries (Grosz & Sidner, 1986).

The overall framework in which the planner functions is shown in Figure 4-1. The dialogue system is continually updating its relational model based on conversational moves by itself and the user, and is dynamically selecting its next utterance on the basis of this model, discourse context, and task goals.

4.1.2 Previous Work in Discourse Planning

The action selection problem (deciding what an autonomous agent should do at any point in time) for conversational agents includes choosing among behaviors with an *interactional function* such as conversation initiation, turn-taking, interruption, feedback, etc., and behaviors with a *propositional function* such as conveying information. Within computational linguistics, the dominant approach to determining appropriate propositional behaviors has been to use a speech-act-based discourse planner to determine the semantic content to be conveyed (also known as “content selection”). Once the content is determined, other processes are typically used to map the semantic representations onto the words the agent actually speaks. Recent text generation systems break this process down into: content ordering (rhetorical organization); sentence planning (aggregating content into sentence-sized units, selecting lexical and syntactic elements); and syntactic realization (Stone & Doran, 1997).

This approach to discourse planning is based on the observation that utterances constitute speech acts (Searle, 1969), such as requesting, informing, wanting and suggesting. In addition, humans plan their actions to achieve various goals, and in the case of communicative actions, these goals include changes to the mental states of listeners. Thus, this approach uses classical “static world” planners (e.g., STRIPS (Fikes & Nilsson, 1971)) to determine a sequence of speech acts that will meet the agent's goals in a given context. One of

the major advantages of plan-based theories of dialog is that language can be treated as a special case of other rational non-communicative behavior.

Conversation to achieve social goals (such as small talk) places many theoretically interesting demands on natural language dialogue systems, a number of which have not been adequately – or at all – addressed by existent approaches to discourse planning. A discourse planner for social dialogue must be able to manage and pursue multiple conversational goals (Tracy & Coupland, 1991), some or all of which may be persistent or non-discrete. For example, in casual small talk, where there are apparently no task goals being pursued, interlocutors are conscious, nevertheless, of multiple goals related to conversation initiation, regulation and maintenance (Cegala et al., 1988). Even in "task-oriented" interactions, speakers may also have several interpersonal goals they are pursuing, such as developing a relationship (e.g., befriending, earning trust) or establishing their reputations or expertise. It is not sufficient that a discourse 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 interpersonal goals of the agent. In addition, many goals, such as intimacy or face goals (Coupland, Coupland, & Robinson, 1992) (Goffman, 1983), 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 (Hanks, 1994). 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.

Discrete reasoning yields a greatly underspecified solution for an agent which must reason about face threats, power, solidarity, and relative goal priorities, especially in social talk in which almost any topic can be raised at any given time, but at varying costs and benefits to the initiating agent.

4.2 REA

The relational dialogue planner was developed for and integrated into an embodied conversational agent named "REA" who performs the role of a real-estate salesperson. Real estate sales was selected as an application domain because a significant portion of the interaction between the agent and a buyer is interpersonal in nature. Within this domain, the initial agent-buyer interview was modeled by studying videotapes of real agent-buyer interactions. In these interviews the real estate agent pursues several goals in parallel, including:

- The task goals of determining clients' housing preferences (size, location, style, etc.) and buying ability.
- The interpersonal goal of establishing trust in the salesperson and reducing clients' fear about such a big purchase.
- The interpersonal goal of establishing the agent's expertise and credibility.

REA is a real-time, multimodal, life-sized ECA who plays the role of a real estate agent, interviewing users about their housing needs and showing them virtual houses that she has for sale (Cassell et al., 1999; Cassell, Bickmore, Vilhjálmsón, & Yan, 2000). Figure 4-2 shows a picture of a user interacting with REA.

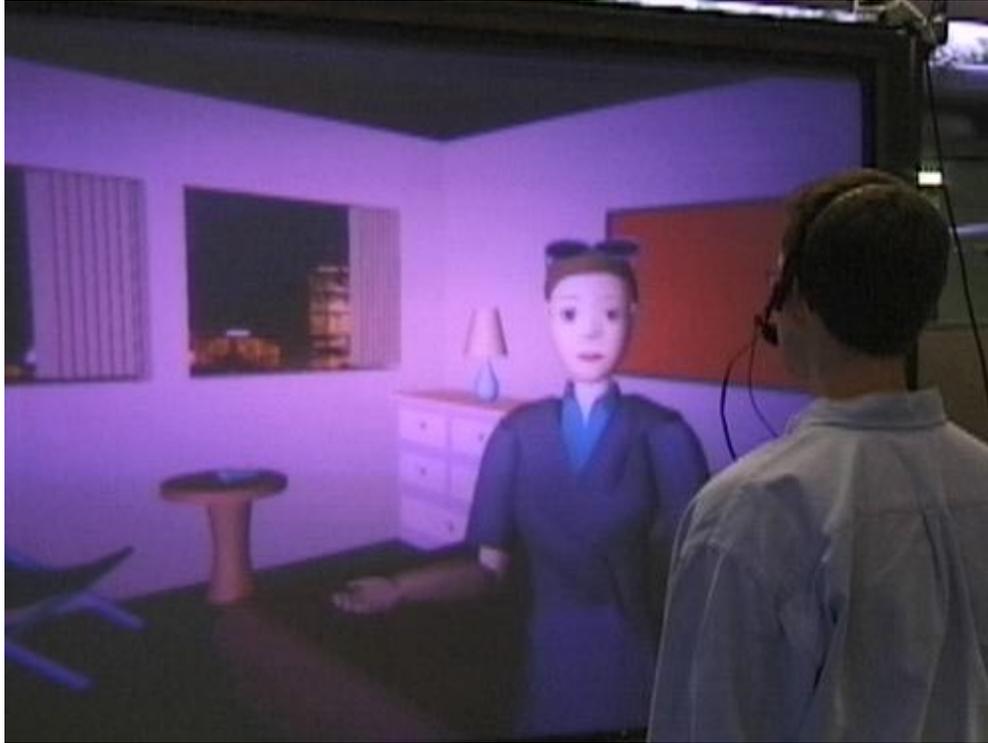


Figure 4-2. User Interacting with Rea

REA 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 gestural output. The system currently consists of a large projection screen on which REA is displayed and which the user stands in front of. Two cameras mounted on top of the projection screen track the user's head and hand positions in space. Users wear a microphone for capturing speech input. A single SGI Octane computer runs the graphics and conversation engine of REA, while several other computers manage the speech recognition and generation and image processing

REA is able to conduct a conversation describing the features of the task domain while also responding to the users' verbal and non-verbal input. When the user makes cues typically associated with turn taking behavior such as gesturing, REA allows herself to be interrupted, and then takes the turn again when she is able. She is able to initiate conversational error correction when she misunderstands what the user says, and can generate combined voice, facial expression and gestural output. A simple discourse model is used for determining which speech acts users are engaging in, and resolving and generating anaphoric references.

4.2.1 REA Architecture

Figure 4-3 shows the modules of the REA architecture. Central to this architecture is the division of conversational contributions into *propositional functions* and *interactional functions*. The interactional discourse functions are responsible for regulating the structure of the interaction between the participants (e.g., turn-taking), while propositional functions shape

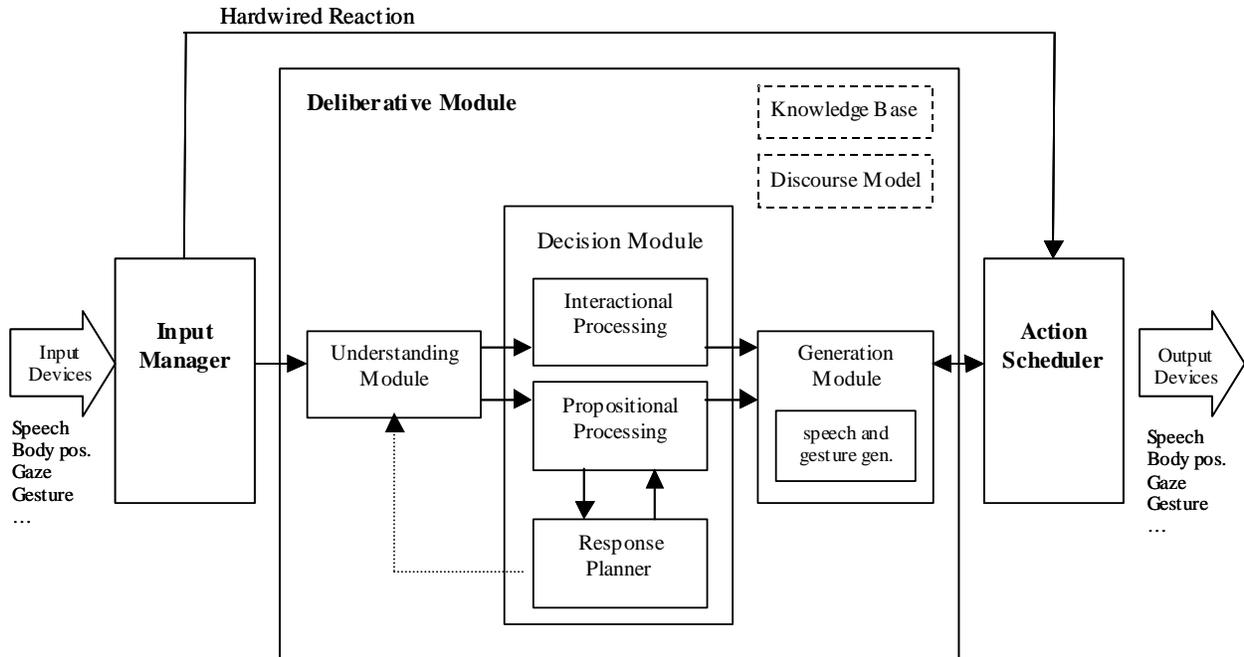


Figure 4-3. Rea Architecture

the actual content communicated. Both functions may be fulfilled by the use of a number of available communication modalities.

In this design, input is accepted from as many modalities as there are input devices. However the different modalities are integrated into a single semantic representation that is passed from module to module. This representation is a KQML frame (Finin et al., 1994) which has slots for interactional and propositional information so that the regulatory and content-oriented contribution of every conversational act can be maintained throughout the system.

The categorization of behaviors in terms of their conversational functions is mirrored by the organization of the architecture which centralizes decisions made in terms of functions (in the deliberative module), and moves to the periphery decisions made in terms of behaviors (the Input Manager and Action Scheduler).

In addition, the Input Manager and Action Scheduler can communicate through a hardwired reaction connection, to respond immediately (under 200 msec.) to user input or system commands. Tracking the user with gaze shifts as they move is an example of a reactive conversational behavior. The other modules are more "deliberative" in nature and perform non-trivial inferencing actions that can take multiple real-time cycles to complete. REA is implemented in C++ and CLIPS, a rule-based expert system language.

4.2.2 Discourse Planner Interface

The discourse planner primarily interfaces with the Reaction Module by providing it with the next utterance that REA should make at any point in the conversation. The planner monitors what is going on in the conversation by analyzing both propositional and

interactional parts of frames representing user inputs and updates its relational model accordingly. It also provides “expectations” to the Input Manager, so that that module is better able to interpret user inputs (e.g., if REA just asked the user how much of a down payment they can make, the Input Manager is instructed to expect a dollar amount, if REA just made a small talk move the Input Manager is instructed to accept any utterance).

4.3 Relational Model

The dimensional model of relationship presented in section 3.9.5 (based on work by Svennevig (Svennevig, 1999)) was used as the basis for representing the relationship between REA and the user. In particular, I was interested in modeling the interaction between language and relationships in a single face-to-face conversation, and this model provided an appropriate amount of granularity without requiring specific beliefs and intentions to be modeled.

Svennevig’s relational model has four dimensions, a refinement of the familiar power and social distance model used frequently in psychology, sociology and sociolinguistics. Power is the ability of one interactant to control the behavior of the other. Svennevig provided significant rationale for decomposing the social distance dimension into three components: solidarity, familiarity and affect.

Solidarity is defined as “like-mindedness” or having similar behavior dispositions (e.g., similar political membership, family, religions, profession, gender, etc.), and is very similar to the notion of social distance used by Brown and Levinson in their theory of politeness (Brown & Levinson, 1978). There is a correlation between frequency of contact and solidarity, but it is not necessarily a causal relation (Brown & Levinson, 1978; Brown & Gilman, 1972).

The third dimension of Svennevig’s relational model is labeled *familiarity*. Based on social penetration theory (Berscheid & Reis, 1998), which claims to account for the establishment and growth of interpersonal relationships, this dimension describes the way in which relationships develop through the reciprocal exchange of information, beginning with relatively non-intimate topics and gradually progressing to more personal and private topics. The growth of a relationship can be represented in both the breadth (number of topics) and depth (public to private) of information disclosed.

The fourth and final dimension of Svennevig's model is *affect*. This represents the degree of liking the interactants have for each other, and there is evidence that this is an independent relational attribute from the above three (Brown & Gilman, 1989). In Pautler's computational model of social perlocutions, affect is the only dimension of relationship modeled (Pautler, 1998).

Elements of this four-dimensional model are used to represent the dynamically-changing relationship between REA and the user. While I currently represent the model as a single entity per dyadic relationship, in fact people frequently form situation-specific representations of others and their relationships with them (they treat their colleagues differently in the gym than in the conference room). Thus, the above model may need to be replicated for different activity types, social roles, or contexts of use (Berscheid & Reis, 1998). The maintenance and integration of such a network of relational models is currently beyond the scope of this work, but provides an interesting area for future research.

4.4 Small Talk

The relational discourse planner has only one degree of freedom in its output, namely the selection of the agent's next utterance. Within the context of achieving task goals, this restricts its relational strategies to appropriate sequencing of task utterances (e.g., deferring threatening topics) and the timely use of social dialogue moves. The social dialogue moves that REA had at her disposal were forms of small talk—questions (“How about this weather?”), statements (“Sorry about my voice, this is some engineer’s idea of natural sounding.”) and stories that do not have direct bearing on her task goals. This section discusses the relational functions of small talk.

It is commonly thought that small talk is what strangers do when they must share a small space for a large period of time, but in general it can be taken as any talk in which interpersonal goals are emphasized and task goals are either non-existent or de-emphasized (including social chit chat, conversational stories, asides). As illustrated above, within task-oriented encounters, small talk can help humans or agents to achieve their goals by "greasing the wheels" of task talk. It can serve a transitional function, providing a ritualized way for people to move into conversation in what may be an otherwise awkward situation (Jaworski & Coupland, 1999). Small talk can also serve an exploratory function by providing a conventional mechanism for people to establish their capabilities and credentials. The realtor in the dialogue cited above, for example, later used small talk to demonstrate her skills by telling a short anecdote about how she had sold a house to her very own tenant, and how successful that sale had been. Small talk can build solidarity if the conversation involves a ritual of showing agreement with and appreciation of the conversational partner's utterances (Malinowski, 1923) (Schneider, 1988) (Cheepen, 1988). People and agents can use small talk to establish expertise, by relating stories of past successful problem-solving behavior, and to obtain information about the other that can be used indirectly to help achieve task goals (e.g., that the client is pregnant increases the probability that the person will require a two-bedroom or larger home). Finally, small talk can be used by partners in long-term relationships to assess each other's moods at the beginning of interactions (Picard, 1997).

Small talk can be used to address the face needs of interlocutors. In small talk, interlocutors take turns showing agreement with and appreciation of the contributions of the speaker, and in so doing enhance each other's face (Cheepen, 1988; Schneider, 1988). This builds solidarity among the interlocutors by demonstrating their "like mindedness". Of course, small talk can also be used in social situations as a prelude to other, more personal kinds of talk (such as "getting acquainted talk" (Svennevig, 1999)), once the interlocutors decide that they want to move on to the next stage of their relationship. Small talk can also be used to address interlocutor's face by defusing awkward silences between strangers, such as in waiting rooms or airplanes (Malinowski, 1923; Schneider, 1988). This is more of a defensive use of small talk, in which the interlocutors are attempting to establish only a minimal level of solidarity.

4.4.1 How Small Talk Works

The topics in small talk are highly constrained, and typically begin with subjects in the interlocutors' immediate shared context (e.g., the weather), since that is both safe and can be presumed to be in the common ground. Topics can then either progress out to the shared sociocultural context (e.g., economy, "light politics"), or in to personal topics of the

participants. The former approach is more typically followed in social contexts (e.g., parties) while the latter is more typical of strangers who must address an awkward silence between them (Schneider, 1987).

When used to address positive face wants, interlocutors show increased attentiveness towards each other. Stylistically, then, small talk can be seen as a kind of ostensible communication (Clark, 1996) in which the interlocutors are pretending to be close friends or acquaintances, while keeping the discourse topics at a safe level of interpersonal distance. This being the case, interlocutors engaged in small talk show signs of positive affect in their speech, conveying some of the signs of "interpersonal warmth," including such behaviors as (Andersen & Guerrero, 1998):

- Proxemic behaviors: close conversational distance, direct body orientation, forward leans, communicating at the same level or in the same physical plane
- Oculestic behaviors: increased gaze, mutual eye contact, decreased eye movements
- Kinesic behaviors: smiling, general facial pleasantness, affirmative head nods, gestural animation, head tilts, bodily relaxation, lack of random movement, open body positions, postural congruence
- Vocalic behaviors: more variation in pitch, amplitude, duration and temp; reinforcing interjections such as "uh-huh" and "mm-hmmm"; greater fluency, warmth, pleasantness, expressiveness, and clarity; smooth turn-taking

Structurally, small talk has been characterized (Schneider, 1988) in terms of an initial question-response pair, followed by one of several types of third moves (echo question, check-back, acknowledgement, confirmation, evaluation), followed by zero or more synchronized "idling" moves. An example of such an exchange reported by Schneider is:

A: It's a nice morning, isn't it?

B: It's very pleasant.

A: It is really, it's very pleasant, yes.

B: Mhm.

Topic introduction also follows a number of structural constraints. Topics are negotiated among the interlocutors, rather than simply introduced by one speaker. The constraints on topic include the following (Svennevig, 1999):

- reportability - a presumption of interest by the interlocutors, established in negotiation. Topics can be proposed via "topical bids" ("I just got back from vacation.") and taken up via "topicalizers" ("Oh yea?") which indicate interest in the topic.
- projectability - a plan for the topic talk should be indicated in the proposal, usually by means of indicating the genre to be used -- narrative, argumentation, exposition, etc. ("oh, that reminds me of a story").
- local connectedness - contributions are fit to the preceding turn.
- progressivity - a topic is continued as long as it progresses (as long as there is new material, or until structural disfluencies occur).
- Interlocutors have a preference for gradual topic transition (Sacks, 1995), and sensitive topics can be broached in a stepwise and collaborative manner which displays the participants' mutual willingness to enter into it, and avoids dispreferred moves and other sudden shifts.

There are also constraints on the introduction of small talk within other types of talk. For example, in conversational frames in which there is an unequal power balance and some level of formality (e.g., job interviews), only the superior may introduce small talk in the medial phase of the encounter (Cheepen, 1988).

Other style constraints include the increased importance of politeness maxims and the decreased importance of Gricean "maximally informative communication" maxims, and the obligatory turn-taking mentioned above (one interlocutor cannot hold floor for the duration of the encounter).

In sum, as illustrated in Figure 4-4, relative to the strategies described above (and the relational dimensions they affect) small talk:

- Avoids *face threat* (and therefore maintains *solidarity*) by keeping conversation at a safe level of depth.
- Establishes *common ground* (and therefore increases *familiarity*) by discussing topics that are clearly in the context of utterance
- Increases *coordination* between the two participants by allowing them to synchronize short units of talk and nonverbal acknowledgement (and therefore leads to increased liking and positive *affect*).
- Requires the demonstration of *reciprocal appreciation* for each other's contributions (and therefore leads to increased *solidarity*).

4.5 Discourse Planner Design

This section presents the technical details of the discourse planner's design. I first discuss the elements of Svennevig's relational model used to represent the user-agent relationship, and how these are continually updated during conversation. I then present the algorithms that

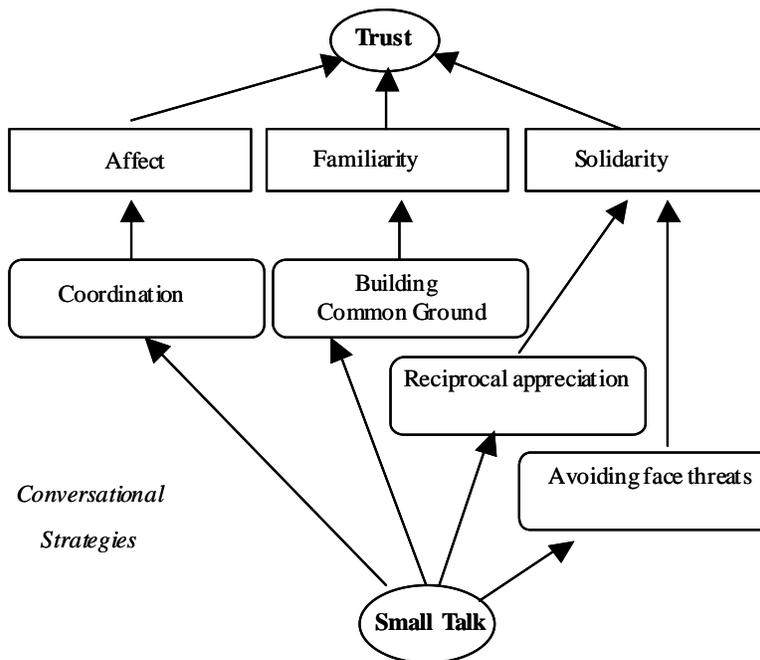


Figure 4-4. How Small Talk Works

are used to select next agent utterances on the basis of this model, discourse context and agent goals.

4.5.1 User Modeling of Human-Computer Relationships

The social relationship between the user and a computer agent represents an underexplored aspect of user modeling. In the implementation of my model, three of Svennevig's five relational dimensions--familiarity/depth, familiarity/breadth, and solidarity (Svennevig, 1999) – are used, each represented as a scalar ranging from zero to one, with increasing values representing increasing closeness. These elements of the user model are updated dynamically during the interaction with the user. In fact, the current implementation does some amount of assessing user state, and adapting to it, but also engages in attempts to *change* user state – choosing behaviors that are intended ultimately to increase the user's trust in the system.

In the implementation, conversational topics are represented as objects that include measures of minimum and maximum 'social penetration' or invasiveness as two of their attributes. For example, talking about the weather does not represent a very invasive topic, whereas talking about finance does. Given these representations, depth of familiarity is updated based on the maximum invasiveness of all topics discussed, and breadth of familiarity is updated based on the number of topics discussed. The model of user solidarity with the system should ultimately be updated based on the similarity of the user's and agent's values and beliefs. However, since solidarity has also been observed to increase based on the number of interactions two individuals have, the current model simply updates solidarity as a function of the number of turns of conversation that the user and agent have engaged in.

More formally, if $T = \{t_1, t_2, \dots, t_J\}$ is the set of possible conversational topics the agent can discuss, $T_H \subseteq T$ is the set of topics already discussed, $T_C \subseteq T$ is the current set of topics under discussion, $D_{MIN}:T \rightarrow 0..1$ and $D_{MAX}:T \rightarrow 0..1$ represent the minimum and maximum social penetration for a topic (depth of familiarity), respectively, N_{moves} is the number of conversational moves made by the agent thus far in the conversation and $M = \{m_1, m_2, \dots, m_K\}$ is the set of possible conversational moves the agent can make, then the relational model is updated as follows.

$$FamiliarityDepth = \frac{Maximum(\{D_{MIN}(i) \mid i \in T_H \cup T_C\})}{Maximum(\{D_{MAX}(j) \mid j \in T\})}$$

$$solidarity = \frac{N_{moves}}{|M|}$$

$$FamiliarityBreadth = \frac{|T_H|}{|T|}$$

One final element of the user model is a set of topics $T_R \subseteq T$ which are relevant to the user throughout the conversation. This set is initialized to topics regarding readily apparent features of the system and the immediate context of utterance -- the setting the user will be in when using the system -- such as REA's appearance and voice, the projection screen and microphone, the lab the system is situated in, MIT, Cambridge, Boston, the weather outside,

etc. This set defines the topics that can readily be discussed with anyone who walks up to the system for the first time, and is thus important for determining topics for small talk, to increase common ground. Currently this set is not updated during use of the system, but ideally it would be expanded as more is learned about the user.

Following the model presented in Chapter 3 for face threat, the degree of threat for any given move m_i is computed as a scalar quantity based on the relational model as follows, given that $A:M \rightarrow 2^T$ is the set of topics a move is "about", $TC: \langle 2^T, 2^T \rangle \rightarrow 0..1$ returns the degree of "topic coherence" between two sets of topics, ranging from 1 if the sets share any common members to 0 if the two sets of topics have nothing in common, $S = \{ 'STORY', 'QUERY', 'STATEMENT' \}$ is the set of possible speech acts the agent can make, and $SA: M \rightarrow S$ provides the class of speech act for a conversational move.

$$\text{face threat}(m_i, \text{familiarity/depth}, \text{familiarity/breadth}, \text{solidarity}) = SP_{\text{threat}} \times G_{SP_{\text{threat}}} + SAI_{\text{threat}} \times G_{SAI_{\text{threat}}} + SAC_{\text{threat}} \times G_{SAC_{\text{threat}}}$$

Where,

$$SP_{\text{threat}} = \text{Maximum}(\{ D_{\text{MIN}}(m_i) - \text{FamiliarityDepth}, 0 \})$$

$$SAI_{\text{threat}} = \text{if } \text{solidarity} \geq S_{\text{MIN}} \text{ then } 0 \text{ else}$$

$$\quad \text{if } SA(m_i) = 'STORY' \text{ then } SA_{\text{STORY}}$$

$$\quad \text{else if } SA(m_i) = 'QUERY' \text{ then } SA_{\text{QUERY}}$$

$$\quad \text{else } SA_{\text{STATEMENT}}$$

$$SAC_{\text{threat}} = 1 - TC(A(m_i), T_C)$$

Here, SP_{threat} is the amount of threat due to violations of social penetration theory, SAI_{threat} is the 'intrinsic' threat of the speech act (conditioned upon the solidarity in the relationship), and SAC_{threat} is the amount of threat due to 'unexpected' topics (those with low coherence to the preceding discourse context).

$G_{SP_{\text{threat}}}$, $G_{SAI_{\text{threat}}}$, $G_{SAC_{\text{threat}}}$ are constant gains (range ≥ 0 , see Table 4-1 for values used in the examples in section 4.6).

SA_{STORY} , SA_{QUERY} , $SA_{\text{STATEMENT}}$ are constants describing the degree of threat from telling a story, asking a question, or making a statement, respectively, if an appropriate level of solidarity (S_{MIN}) has not been established.

	Normal REA	Goal-oriented REA	Chatty REA
G_{ENABLED}	0.5	1.0	0.2
G_{GOAL}	0.7	1.0	0.2
G_{Spthreat}	5.0	1.0	5.0
G_{SPENABLE}	1.0	1.0	1.0
$G_{\text{TOPICENABLE}}$	1.0	0.3	0.6
$G_{\text{SACthreat}}$	2.2	1.0	2.2
$G_{\text{FACETHREAT}}$	-1.5	0.0	-1.5
$G_{\text{SAIthreat}}$	1.0	1.0	1.0
$G_{\text{RELEVANCE}}$	0.5	0.0	0.5

Table 4-1. Gain Values Used in Examples (see Section 4.6)

4.5.2 Discourse Planning for Mixed Task and Social Dialog

For the purpose of trust elicitation and small talk, I have constructed a new kind of discourse planner that can interleave small talk and task talk during the initial buyer interview, based on the model outlined above. Given the requirements to work towards the achievement of multiple, non-discrete goals in a dynamically changing environment, I have moved away from static world discourse planning, and am using an activation network-based approach based on Maes' *Do the Right Thing* architecture (Maes, 1989). In the implementation of this architecture, each node in the network represents a conversational move that REA can make, and links between the nodes represent various enabling and disabling conditions which hold among the moves (e.g., talking about the Boston real estate market introduces the topic of real estate thereby making it easier to introduce other real estate topics; see Figure 4-5).

Planning takes place as a spreading activation process that uses information from the current state of the conversation and relational model to determine which moves are more likely to succeed, along with information from the task goals so that REA can prioritize her possible moves to ensure that she addresses goals based on their relative importance, and adapted to her model of the user. Plans, ordered sequences of conversational moves, are represented as paths in the network, and are determined via the spreading activation mechanism. Although planning is performed before a move is selected, the mechanism does not return a complete plan like classical planners do. Instead, moves are selected for execution as necessary to achieve the unsatisfied task goals. During the activation process, energy is moved backward from the task goals to moves which directly lead to their achievement, from there to moves which enable those moves, and so on. Forward propagation is simultaneously performed by flowing energy into moves which can immediately be performed given the conversational state and relational model, and from there to moves which are enabled by those moves, and so on. The resulting paths in the network with the highest activation energy are thus those that are executable and best lead to satisfaction of the task goals.

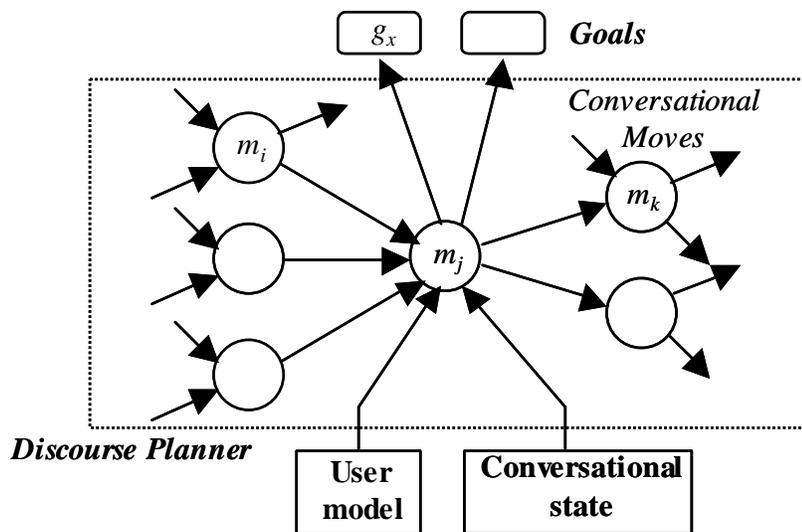


Figure 4-5. Conversational Moves in the Activation Network

Since the role of goals is to simply inject energy into moves which lead to their achievement, the network can straightforwardly be extended to work towards the achievement of non-discrete goals by simply varying the amount of energy the goals provide based not only on their relative priority, but on the difference between their current and desired degrees of satisfaction (the larger the discrepancy the more energy they provide). The pursuit of multiple goals can also be handled in parallel, with goal conflicts handled by inconsistent moves sending negative activation energy to each other. In addition, if a move does not produce an expected result, or if the conversational or relational states change in unanticipated ways, rather than re-planning (as a classical planner would do), the next best network path is automatically selected for execution. Thus, this architecture supports the requirements for planning to achieve multiple, non-discrete goals in a dynamic environment.

In addition, this architecture provides many features that are of particular use to designers of conversational agents. A discourse planner modeled in this manner provides enormous flexibility in designing agents whose conversational behavior vary by degree of goal-directedness, politeness, coherence, relevance or even deliberation (vs. opportunism), simply by changing the appropriate numeric gains controlling the amount of activation energy propagated under certain conditions. Since the spreading activation process incrementally estimates the best path to take through the network, it represents a form of "anytime" planner that can be stopped at any time to provide the best action to execute, although the longer it is allowed to run the better the result (Drummond & Bresina, 1990). Thus, the architecture provides the capability to transition smoothly from deliberative, planned behavior to opportunistic, reactive behavior by varying the length of time the activation energy propagation algorithm runs.

Activation Network Architecture

I adapted Maes' architecture for discourse planning in order to support mixed task and social dialog in REA, to adapt to the model of the user, and to dynamically attempt to change the user's relationship with the system. During task talk, REA asks questions about users' buying preferences, such as the number of bedrooms they need. During small talk, REA can talk about the weather, events and objects in her shared physical context with the user (e.g., the lab setting), or she can tell stories about the lab, herself, or real estate.

REA's contributions to the conversation are planned in order to minimize the face threat to the user while pursuing her task goals in the most efficient manner possible. That is, REA attempts to determine the face threat of her next conversational move, assesses the solidarity and familiarity which she currently holds with the user, and judges which topics will seem most relevant and least intrusive to users. As a function of these factors, REA chooses whether or not to engage in small talk, and what kind of small talk to choose. The selection of which move should be pursued by REA at any given time is thus a non-discrete function of the following factors:

- From Maes:
 1. Task goals -- REA has a list of prioritized goals to find out about the user's housing needs in the initial interview. Conversational moves that directly work towards satisfying these goals (such as asking interview questions) are preferred.

2. Logical preconditions -- Conversational moves have logical preconditions (e.g., it makes no sense for REA to ask users what their major is until she has established that they are students), and are not selected for execution until all of their preconditions are satisfied. Moves whose preconditions are satisfied by the user model and conversational state are given an increment of "forward chaining" energy. The move A which, when selected, will enable another move B, passes some of its activation energy forward from A to B. In addition, a move A which helps satisfy the preconditions of move B causes activation energy to flow from B to A, providing a "backward chaining" flow of energy.
- To deal with face threat avoidance:
 3. Face threat -- Moves that are expected to cause face threats to the user are dis-preferred, including face threats due to social penetration theory violations, speech act type or topic incoherence.
 4. Familiarity/Depth enablement -- In order to avoid face threats due to social penetration theory violations, REA can plan to perform small talk in order to "grease the tracks" for task talk, especially about sensitive topics like finance. To support this, energy is passed from moves whose familiarity/depth preconditions are not satisfied to those moves which would satisfy these preconditions if selected.
 - To deal with topic management:
 5. Relevance -- Moves that involve topics in the list of topics known to be relevant to the user are preferred.
 6. Topic enablement -- REA can plan to execute a sequence of moves that gradually transition the topic from its current state to one that REA wants to talk about (e.g., from talk about the weather, to talk about Boston weather, to talk about Boston real estate). Thus, energy is propagated from moves whose topics are not currently active to moves whose topics would cause them to become current.

More formally, given the set of agent goals $G = \{g_1, g_2, \dots\}$, the degree of satisfaction, $S_G: G \rightarrow 0..1$, and priority, $P_G: G \rightarrow 0..1$, for each goal, each move is assigned the following activation energy during each update cycle.

$$\alpha_i^0 = 0$$

$$\alpha_i^t = \text{decay}(\alpha_i^{t-1}) +$$

$$E_{GOAL}(m_i) * G_{GOAL} + E_{ENABLED}(m_i) * G_{ENABLED} + E_{FORWARD}(m_i) * G_{FORWARD} + E_{BACKWARD}(m_i) * G_{BACKWARD} +$$

$$E_{RELEVANCE}(m_i) * G_{RELEVANCE} + E_{FACETHREAT}(m_i) * G_{FACETHREAT} +$$

$$E_{TOPICENABLE}(m_i) * G_{TOPICENABLE} + E_{SPENABLE}(m_i) * G_{SPENABLE}$$

Where,

G_{GOAL} , $G_{ENABLED}$, $G_{RELEVANCE}$, $G_{TOPICENABLE}$, $G_{SPENABLE}$, $G_{FORWARD}$, and $G_{BACKWARD}$ are gain constants (≥ 0), and $G_{FACETHREAT}$ is a negative gain constant (≤ 0). Modification of these gains allows the agent to be made more or less goal-oriented (by changing G_{GOAL}), more or less polite (by changing $G_{FACETHREAT}$) or more or less deliberate in how topics are advanced (by changing $G_{TOPICENABLE}$). Table 4-1 gives the gain values used in the examples presented in section 4.6.

$$decay(\alpha_i^t) = \frac{|M| \times \pi}{\sum_{j \in M} \alpha_j^t}$$

π is a normalization constant which controls the total amount of energy available in the network (the 'mean level of activation').

$$E_{GOAL}(m_i) = \sum_{g \in C_G(m_i)} (1 - S_G(g_i)) * P_G(g_i)$$

$C_G: M \rightarrow 2^G$ is the set of goals that a move directly contributes to the satisfaction of.

$E_{ENABLED}(m_i) = 1$ if all logical preconditions of the move are satisfied, 0 otherwise.

$$E_{FORWARD}(m_i) = \sum_{m_j \in M_{ENABLES}(m_i)} \alpha_j^{t-1}$$

$$E_{BACKWARD}(m_i) = \sum_{m_k \in M_{ENABLEDBY}(m_i)} \alpha_k^{t-1}$$

$M_{ENABLES}: M \rightarrow 2^M$ is the set of moves which have at least one logical precondition directly satisfied through the execution of a given move, and $M_{ENABLEDBY}: M \rightarrow 2^M$ is the inverse (the set of moves which, when executed, satisfy at least one logical precondition of the given move).

$$E_{RELEVANCE}(m_i) = TC(A(m_i), T_R)$$

$$E_{FACETHREAT}(m_i) = \text{facethreat}(m_i, \text{FamiliarityDepth}, \text{FamiliarityBreadth}, \text{solidarity})$$

$$E_{TOPICENABLE}(m_i) = \sum_{m_j \in M | A(m_j) - A(m_i) \neq \{ \}} TC(A(m_i), A(m_j)) \times \alpha_j^{t-1}$$

Here, one conversational move 'topic enables' another when it can introduce additional topics ($m_j \in M | A(m_j) - A(m_i) \neq \{ \}$) that are coherent to some degree with the move in question,

thus allowing the move in question to be introduced by way of first talking about the additional topics.

$$E_{SPENABLE}(mi) = \sum_{mj \in SPENABLE(mi)} \alpha_j^{t-1}$$

$$SPENABLE(mi) = \{mj \mid mj \in M \wedge D_{MOVEMIN}(mj) \leq D_{MOVEMIN}(mi) \wedge \\ D_{MOVEMIN}(mj) \leq D_{MOVEMAX}(mi) \wedge \\ FamiliarityDepth \geq D_{MOVEMIN}(mi)\}$$

$$D_{MOVEMAX}(mi) = \text{Maximum}(\{D_{MAX}(x) \mid x \in A(mi)\})$$

$$D_{MOVEMIN}(mi) = \text{Minimum}(\{D_{MIN}(x) \mid x \in A(mi)\})$$

This last factor propagates energy from a move which is currently dis-preferred because of a social penetration theory violation to moves which could enable it by increasing *FamiliarityDepth* when executed. One move (mi) enables another (mj) in this way when mi has lower intimacy requirements than mj, when the two moves overlap in their intimacy ranges, and when mi can be readily used given the current level of intimacy with the user.

Given the above activation energy update rule, a threshold of activation, θ , and a threshold decrement, $0 < \theta_{DECREMENT} < 1$, planning in the network proceeds as follows.

$\theta \leftarrow \theta_{INITIAL}$

while a move has not been selected do

compute α_i for all moves

select move m_i with maximum α_i such that $\alpha_i > \theta$ and $E_{ENABLED}(m_i) = 1$

*if no such move is found then $\theta \leftarrow \theta * \theta_{DECREMENT}$*

In the current implementation, the dialogue is entirely REA-initiated, and user responses are recognized via a speaker-independent, grammar-based, continuous speech recognizer (currently IBM ViaVoice). The active grammar fragment is specified by the current conversational move, and for responses to many REA small talk moves the content of the user's speech is ignored; only the fact that the person responded at all is enough to advance the dialogue. This strategy may seem to indicate the opposite of user modeling but, in practice, much human-human small talk proceeds along similar lines and as described above, the tight temporal coordination of units is actually more important than content.

At each step in the conversation in which REA has the floor (as tracked by a conversational state machine in REA's Reaction Module (Cassell et al., 2000)), the discourse planner is consulted for the next conversational move to initiate. At this point, activation values are incrementally propagated through the network (following the algorithm above) until a move is selected whose preconditions are satisfied and whose activation value is above the specified threshold. Moves are executed differently depending on their type. Task queries consist of REA question/user replay pairs; task and small talk statements consist of a REA statement turn only; and small talk stories and queries consist of a REA contribution/optional user response/REA idle response triples.

Shifts between small talk moves and task moves are marked by conventional contextualization cues--discourse markers and beat gestures. Discourse markers include "so"

on the first small talk to task talk transition, "anyway" on resumption of task talk from small talk, and "you know" on transition to small talk from task talk (Clark, 1996).

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 which can facilitate this. The activation energy from the user relevance condition described above leads to REA starting small talk with topics that are known to be in the shared environment with the user (e.g., talk about the weather or the lab).

Note that this implementation is a simplification of Maes' architecture in that it currently assumes information in the conversational state is monotonic, thus goal protection and action conflicts are not currently dealt with (for example once a given level of familiarity is achieved it is assumed that there are no actions the agent can take that would reduce it). I also assume that each conversational move can only be used once in a given interaction and thus disable moves that have been executed by effectively removing them from the network. Finally, given that the threshold of activation, θ , is decreased on each update cycle, $\theta_{DECREMENT}$ controls the amount of deliberation the network performs by controlling the number of update cycles executed before a move is selected. As long as $\theta_{DECREMENT} < 1$ the algorithm will eventually yield a result unless there are no moves available whose logical preconditions are satisfied. In practice, a $\theta_{DECREMENT}$ of 0.1 (as used by Maes) along with $\theta_{INITIAL}=3\pi$ and $\pi=1/|M|$ results in move selection after just a few update cycles.

Related Work in Activation Network-Based Planning

Goetz recast Maes' networks as connectionist Hopfield networks which perform pattern recognition. In the process he discovered several interesting constraints and shortcomings in Maes' networks, but most importantly demonstrated that if certain non-linearities are added to the update rules the behavior of the network became more stable with respect to persistently pursuing a given plan (Goetz, 1997).

A more recent, probabilistic, reformulation of this approach to planning was taken by Bagchi, et al. (Bagchi, Biswas, & Kawamura, 1996), in which the network consists of actions and explicitly represented propositions which are pre- and post-conditions of the actions. In this architecture, the activation values associated with propositions reflect the probability of their being true, while the values associated with actions reflect their expected utility. The process of spreading activation is used to incrementally estimate these probabilities and utilities using calculations local to each node in the network. In this approach, the action with the highest utility is selected at the end of each update cycle for execution. I did not adopt this probabilistic approach since it has not been extended to deal with non-discrete goals or propositions yet. However, it is a promising direction for future work.

4.6 Example Interactions

In the real estate domain an agent has several task goals--such as finding out information about the user's desired location, price range, house size, and amenities--with varying priorities (price and location are most important). The interaction of these goals with the dynamically changing user model yields a fairly natural conversational behavior for this domain. With minimal tuning of the network gains REA can be made very goal-oriented or very chatty, although finding desired in-between behaviors can require some tuning effort. As

long as $G_{SAC_{threat}}$ is kept high (maintaining coherence) and $G_{RELEVANCE}$ is kept above zero (maintaining some user relevance) the resulting conversational behavior is natural and believable.

There are some limitations of this approach with respect to other forms of planning, however. In the current model the moves in the network represent 'ground level' actions rather than abstract schemata, limiting the flexibility and scalability of the approach relative to classical hierarchical planners. There are also no guarantees of correctness or completeness of the plans produced; the spreading activation approach is a heuristic one. Finally, it is unclear how activation network based planners could deal with the very complex goal interactions or temporal constraints that many classical planners have been designed to handle.

The following examples display some actual output from REA in conversation with a user (user responses are only shown in positions in which they affect the selection of subsequent joint projects). The first example illustrates REA engaging in baseline small talk (Figure 4-6).

In this example, REA opens with small talk moves regarding things she knows she can discuss with the user (1-5). After a few turns, enough familiarity has been established (simply by doing small talk) that REA can move into task talk (6-9), aided by the transition move in turn 5 (linking Boston to real estate). However, before bringing up the topic of finance -- a topic that is potentially very face threatening for the user -- REA decides that additional familiarity needs to be established, and moves back into small talk (10-15). The move at turn 15 transitions the topic from the Media Lab back to real estate, enabling REA to then bring up the issue of how much the user is able to afford (16-17).

In the next example (Figure 4-7), REA does not perform any small talk and sequences the task questions in strictly decreasing order of priority.

	Move	Fam/D	Fam/B	Solidarity
1.	So, where would you like to live?	0.00	0.00	0.00
2.	What can you afford?			
3.	What kind of down payment can you make?			
4.	How many bedrooms do you need?			
5.	Do you need access to the subway?			
6.	Is one bath enough?			
7.	Let me see what I have available.	0.50	0.19	0.20

Figure 4-7. "Task REA"

As a final example, the importance of task goal satisfaction can be reduced (relative to the baseline case) yielding the following dialogue in which REA rambles on at length before getting down to business (Figure 4-8). When task goals become less important, topic coherence and other sources of activation energy become the primary factors in deciding what REA says next.

	Move	Fam/D	Fam/B	Solidarity
1.	How about this weather?	0.00	0.00	0.00
2.	I think winters in Boston are awful.			
3.	How do you like Boston?			
4.	I have lived in Boston all my life. Come to think of it, I have lived inside this room all of my life. It is so depressing.			
5.	Boston is certainly more expensive than it used to be.	0.50	0.19	0.17
6.	So, Where would you like to live?			
7.	How many bedrooms do you need?			
8.	Do you need access to the subway?			
9.	Is one bath enough?	0.60	0.29	0.30
10.	You know, I keep showing the researchers here the same houses, over and over again. Maybe one day I will get lucky.			
11.	Have you been in the Media Lab before?			
12.	Do you know that the Media Lab is going to expand into another building. Things are really going well for the researchers here.			
13.	It is pretty cool do you think?			
14.	They are doing some crazy things in here.			
15.	I have shown houses to lots of students and faculty from M I T. But I always enjoy talking to them.	0.70	0.38	0.50
16.	Anyway, What can you afford?			
17.	What kind of down payment can you make?			
18.	Let me see what I have available.	0.90	0.43	0.57

Figure 4-6. "Social Language REA"

	Move	Fam/D	Fam/B	Solidarity
1.	How about this weather?	0.00	0.00	0.00
2.	I think winters in Boston are awful.			
3.	How do you like Boston?			
4.	Boston is certainly more expensive than it used to be.	0.40	0.14	0.13
5.	So, where would you like to live?			
6.	How many bedrooms do you need?	0.50	0.24	0.20
7.	You know, I keep showing the researchers here the same houses, over and over again. Maybe one day I will get lucky.			
8.	Have you been in the Media Lab before?			
9.	Do you know that the Media Lab is going to expand into another building. Things are really going well for the researchers here.			
10.	It is pretty cool do you think?			
11.	They are doing some crazy things in here.			
12.	I have shown houses to lots of students and faculty from MIT. But I always enjoy talking to them.			
13.	I really love it here at MIT. Do you?			
14.	I suppose you have heard about all of the hacks at MIT? My favorite was when they dressed the dome up to look like a robot.			
15.	It is pretty intense, do you think?			
16.	Are you tired of demos yet?			
17.	That microphone is terrible. I hate using those things.			
18.	Sorry about my voice. This is some engineers idea of natural sounding.	0.60	0.43	0.60
19.	Anyway, what can you afford?			
20.	What kind of down payment can you make?			
21.	Do you need access to the subway?			
22.	Is one bath enough?			
23.	Let me see what I have available.	0.90	0.48	0.73

Figure 4-8. "Chatty REA"

4.7 Summary

The original goals of developing a discourse planner capable of working towards multiple, non-discrete goals in a dynamic environment have been satisfied by the model and implementation presented, and it meets the needs of discourse planning for mixed task and social dialog to assess and adapt to user relational state. An evaluation of the naturalness and ability of this model to build trust in users is described in the following chapter.

Relative to the full space of planning relational dialogue, this model has obviously just scratched the surface. The dimensional model used was a rough first approximation to the belief and intention-based model of relationships that ultimately needs to be used to represent all the nuances of the relationship between an agent and a user. The assessment of the state of this model from verbal and nonverbal behaviors in conversation (and from other sources of evidence), needs to be developed in much greater detail, in addition to assessment of the user's perception of the relationship (important for assessing cooperation, as discussed in Chapter 3). At the strategic, content planning level, there are many additional relational strategies that can be used in conversation in addition to small talk, including meta-relational conversation, continuity behaviors, and talking about the past and future together, that are known to play important roles in many kinds of long-term relationships. Finally, at the text generation level, relational status and goals have a pervasive effect on language form that has only been partly addressed by researchers in computational linguistics. Altogether, this represents an exciting and challenging direction for future work.

Chapter 5

EVALUATION OF FACE-TO-FACE INTERACTION WITH A RELATIONAL AGENT

Will people actually engage in relational activities with a software agent? Does social dialogue have the same psychological effects on people in human-computer interaction as in human-human interaction? This chapter provides some answers to these questions, by presenting the results of a study of people interacting with a life-sized animated real estate agent who used social dialogue in her initial interviews with them about their housing needs. This work is reported in (Bickmore & Cassell, to appear).

5.1 Introduction

In order to evaluate whether an ECA's social language can actually build trust, solidarity, and interpersonal closeness with users, I conducted an empirical study in which subjects were interviewed by REA about their housing needs, shown two "virtual" apartments, and then asked to submit a bid on one of them¹. For the purpose of the experiment, REA was controlled by a human wizard, following scripts identical to the output of the discourse planner described in Chapter 4, but not dependent on automatic speech recognition or computational vision (Dahlback, Jonsson, & Ahrenberg, 1998). The study was a between subjects design with subjects randomly assigned either to a version of REA which used only task-oriented dialogue (TASK condition) or to an identical version which also included the social dialogue (SOCIAL condition). In addition, to evaluate the effects of agent nonverbal conversational behavior on user attitudes and behavior, I also varied whether REA appeared in her usual form, as a life-sized animated character (EMBODIED condition) or as a disembodied voice on the phone (PHONE condition). Together these variables provided a 2x2 experimental design: SOCIAL vs. TASK and EMBODIED vs. PHONE.

The hypotheses for this empirical evaluation follow from the literature on small talk and on interpersonal relations among humans. As discussed in chapter 3, trust is one of the primary outcomes of relational behavior, so I expected subjects in the SOCIAL condition to trust REA more. I also expected them to feel closer to REA, like her more, and feel that they understand her and were understood by her more than in the TASK condition. I expected users to think the interaction was more natural, lifelike, and comfortable in the SOCIAL condition. I also expected users to be willing to pay REA more for an apartment in the SOCIAL condition, given the hypothesized increase in trust. Finally, I expected all of these SOCIAL effects to be amplified given the presence of REA's nonverbal behavior (in EMBODIED condition) relative to the PHONE-only condition.

¹ I designed, implemented, ran and analyzed data from the initial small talk study with REA in EMBODIED condition. Following this, Jennifer Smith and Elisabeth Sylvan designed an extension to the study to test the PHONE condition and ran the experiment with several subjects, which I then completed and analyzed the data for.

5.2 Experimental Method

This was a multivariate, multiple-factor, between-subjects experimental design, involving 58 subjects (69% male and 31% female).

5.2.1 Apparatus

One wall of the experiment room was a rear-projection screen. In the EMBODIED condition REA appeared life-sized on the screen, in front of the 3D virtual apartments she showed (see Figure 5-1), and her synthetic voice was played through two speakers on the floor in front of the screen. In the PHONE condition only the 3D virtual apartments were displayed and subjects interacted with REA over an ordinary telephone placed on a table in front of the screen.

For the purpose of this experiment, REA was controlled via a wizard-of-oz setup on another computer positioned behind the projection screen (see Figure 5-2). The interaction script included verbal and nonverbal behavior specifications for REA (e.g., gesture and gaze commands as well as speech), and embedded commands describing when different rooms in the virtual apartments should be shown. Three pieces of information obtained from the user during the interview were entered into the control system by the wizard: the city the subject wanted to live in; the number of bedrooms s/he wanted; and how much s/he was willing to spend. The first apartment shown was in the specified city, but had twice as many bedrooms as the subject requested and cost twice as much as s/he could afford (they were also told the price was "firm"). The second apartment shown was in the specified city, had the exact number of bedrooms requested, but cost 50% more than the subject could afford (but this time, the subject was told that the price was "negotiable"). The scripts for the TASK and



Figure 5-1. Rea Showing an Apartment

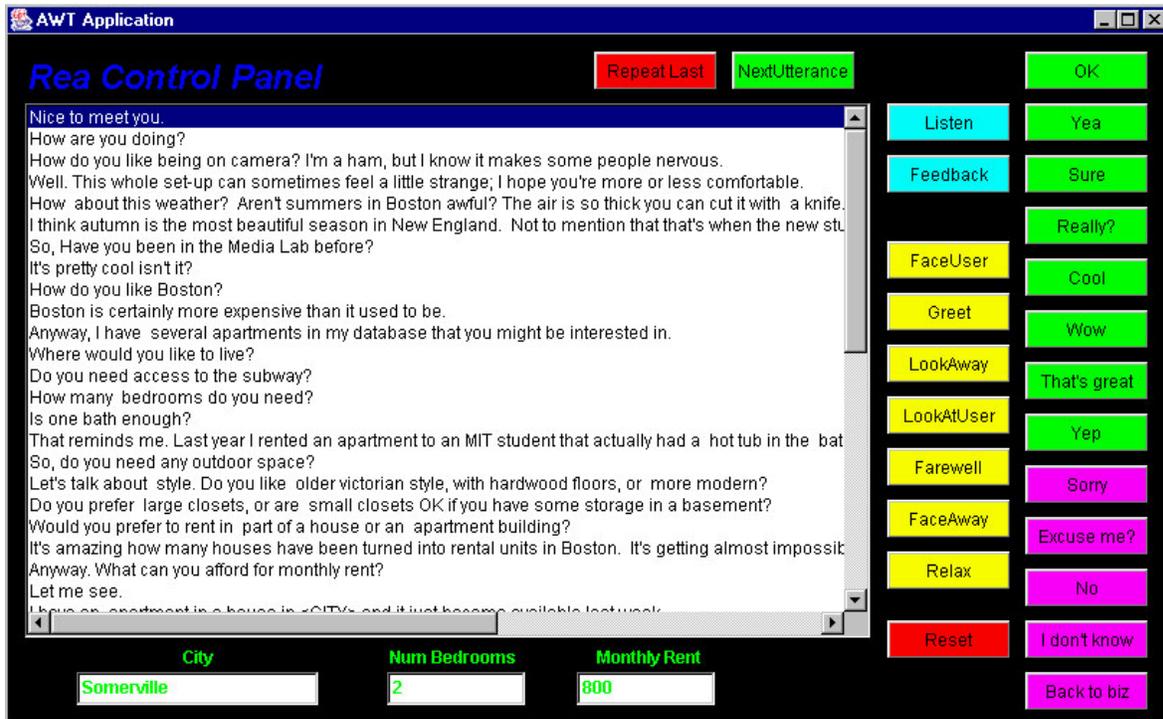


Figure 5-2. Wizard Control Panel for Rea

SOCIAL conditions were identical, except that the SOCIAL script had additional small talk utterances added to it, as described in Chapter 4 (the TASK script had 32 agent speaking turns while the SOCIAL script had 45). The part of the script governing the dialogue from the showing of the second apartment through the end of the interaction was identical in both conditions.

5.2.2 Procedure

Subjects were told that they would be interacting with REA, who played the role of a real estate agent and could show them apartments she had for rent. They were told that they were to play the role of someone looking for an apartment in the Boston area. In both conditions subjects were told that they could talk to REA "just like you would to another person".

Subjects were then shown a brief (one minute) video of REA on a small monitor, giving additional instructions regarding her speech recognition software. The purpose of this was to both reduce the "novelty effect" when REA first appeared on the big projection screen, and to ensure the deception (use of a wizard) was effective. Subjects then interacted with REA, after which they were asked to fill out a questionnaire.

5.2.3 Manipulation check

Three questions concerning the amount of small talk used by REA were included on the questionnaire, both for development feedback and for manipulation checks. That is, subjects were asked, for example, how quickly REA got down to business. If there is a perceivable difference between the small talk and task-only conditions, then subjects should believe that task-only REA got down to business more quickly. All three manipulation check variables

were highly significant. For example, there was a significant difference ($F= 11.2$; $p< .002$) such that users believed that REA got down to business more quickly in the task-only condition than in the small talk condition.

5.2.4 Measures

Subjective evaluations of REA -- including how friendly, credible, lifelike, warm, competent, reliable, efficient, informed, knowledgeable and intelligent she was -- were measured by single items on nine-point Likert scales. Evaluations of the interaction--including how tedious, involving, enjoyable, natural, satisfying, fun, engaging, comfortable and successful it was--were also measured on nine-point Likert scales. Evaluation of how well subjects felt they knew REA, how well she knew and understood them and how close they felt to her were measured in the same manner.

Liking of REA was an index composed of three items--how likeable and pleasant REA was and how much subjects liked her--measured items on nine-point Likert scales. Cronbach's alpha was used as a measure of the internal consistency of this scale--the degree to which these three questions measured the same underlying attitude--with a resulting value of 0.87.

Amount Willing to Pay was computed as follows. During the interview, REA asked subjects how much they were able to pay for an apartment; subjects' responses were entered as \$X per month. REA then offered the second apartment for \$Y (where $Y = 1.5 X$), and mentioned that the price was negotiable. On the questionnaire, subjects were asked how much they would be willing to pay for the second apartment, and this was encoded as Z. The task measure used was $(Z - X) / (Y - X)$, which varies from 0% if the user did not budge from their original requested price, to 100% if they offered the full asking price.

Trust was measured by a standardized trust scale (Wheless & Grotz, 1977) ($\alpha = .93$).

Given literature on the relationship between user personality and preference for computer behavior, I was concerned that subjects might respond differentially based on predisposition. Thus, I also included composite measures for introversion and extroversion on the questionnaire.

Extrovertedness was an index composed of seven Wiggins (Wiggins, 1979) extrovert adjective items: Cheerful, Enthusiastic, Extroverted, Jovial, Outgoing, and Perky. It was used for assessment of the subject's personality ($\alpha = .87$).

Introvertedness was an index composed of seven Wiggins (Wiggins, 1979) introvert adjective items: Bashful, Introverted, Inward, Shy, Undemonstrative, Unrevealing, and Unsparkling. It was used for assessment of the subject's personality ($\alpha = .84$).

Observation of the videotaped data made it clear that some subjects took the initiative in the conversation, while others allowed REA to lead. Unfortunately, the wizard-of-oz script that REA was operating on was not designed to deal with user-initiated talk, and so user initiative often led to REA interrupting the speaker. To assess the effect of this phenomenon, I divided subjects into PASSIVE (below the mean on number of user-initiated utterances) and ACTIVE (above the mean on number of user-initiated utterances). These measures turned out to be independent of introversion/extroversion (Pearson $r=0.042$), and to not be predicted by these latter variables.

5.3 Results

Full factorial single measure ANOVAs were run, with SOCIALITY (Task vs. Social), PERSONALITY OF SUBJECT (Introvert vs. Extrovert), MEDIUM (Phone vs. Embodied) and INITIATION (Active vs. Passive) as independent variables.

5.3.1 Subjective Assessments of REA

Main Effects

In looking at the questionnaire data subjects seemed to feel more comfortable interacting with REA over the phone than face-to-face. Thus, subjects in the phone condition felt that they knew REA better ($F=5.02$; $p<.05$), liked her more ($F=4.70$; $p<.05$), felt closer to her ($F=13.37$; $p<.001$), felt more comfortable with the interaction ($F=3.59$; $p<.07$), and thought REA was more friendly ($F=8.65$; $p<.005$), warm ($F=6.72$; $p<.05$), informed ($F=5.73$; $p<.05$), and knowledgeable ($F=3.86$; $p<.06$) than those in the embodied condition.

Interactions

Subjects felt that REA knew them ($F=3.95$; $p<.06$) and understood them ($F=7.13$; $p<.05$) better when she used task-only dialogue face-to-face; these trends were reversed for phone-based interactions. Task-only dialogue was more fun ($F=3.36$; $p<.08$) and less tedious ($F=8.77$; $p<.005$; see Figure 5-3) when embodied, while social dialogue was more fun and less tedious on the phone. That is, in the face-to-face condition, subjects preferred REA to simply “get down to business.”

These results indicate that REA’s nonverbal behavior was especially inappropriate for social dialogue. REA’s smiles are limited to those related to the ends of turns, and she did not have a model of immediacy or other nonverbal cues for liking and warmth typical of social interaction (Argyle, 1988). This may explain why subjects preferred task interactions face-to-face, while on the phone REA’s social dialogue had its intended effect of making subjects feel that they knew REA better, that she understood them better, and that the experience was more fun and less tedious.

There was a three-way interaction between SOCIALITY, PERSONALITY and MEDIUM ($F=3.96$; $p<.06$) that indicated that extroverts trusted REA more when she used social dialogue in embodied interactions, but there was essentially no effect of user’s personality and social dialogue on trust in phone interactions (see Figure 5-4). Further analysis of the data indicated that this result derived from the substantial difference between introverts and extroverts in the face-to-face task-only condition. Introverts trusted her significantly more in the face-to-face task-only condition than in the other conditions ($p<.03$),

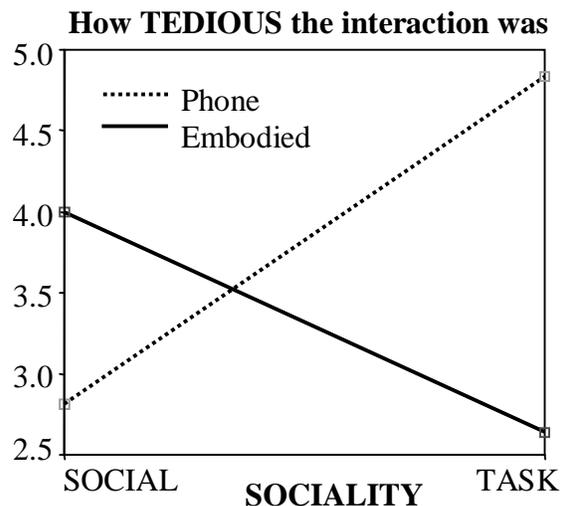


Figure 5-3. Ratings of How TEDIOUS the Interaction was

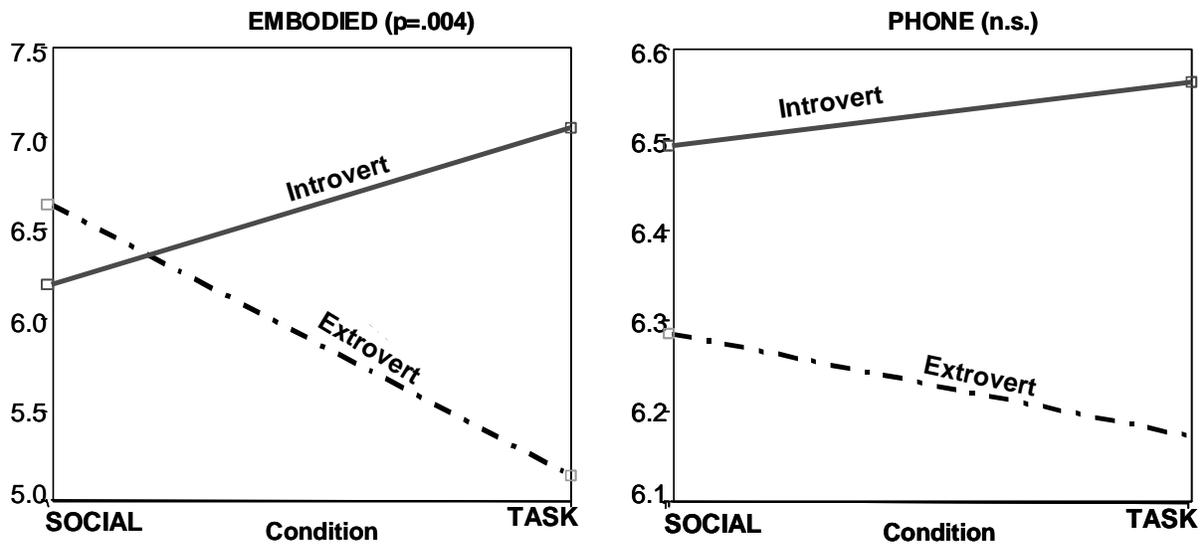


Figure 5-4. Trust in REA

while extroverts trusted her significantly less in this condition than in the other conditions ($p < .01$).

This analysis indicates that the effects on trust may be due to the attraction of a computer displaying similar personality characteristics, rather than the process of trust-building. In the face-to-face, task-only condition both verbal and nonverbal channels were clearly indicating that REA was an introvert (also supported by the comments that REA's gaze-away behavior was too frequent, an indication of introversion (Wilson, 1977)), and in this condition the introverts trust more, and extroverts trust less. In all other conditions, the personality cues are either conflicting (a mismatch between verbal and nonverbal behavior has been demonstrated to be disconcerting to users (Nass, Isbister, & Lee, 2000)) or only one channel of cues is available (i.e. on the phone), yielding trust ratings that are close to the overall mean.

There was, nevertheless, a preference by extroverts for social dialogue as demonstrated by the fact that, overall, extroverts liked REA more when she used social dialogue, while introverts liked her more when she only talked about the task ($F=8.09$; $p < .01$).

Passive subjects felt more comfortable interacting with REA than active subjects did, regardless of whether the interaction was face-to-face or on the phone, or whether REA used social dialogue or not. Passive subjects said that they enjoyed the interaction more ($F=4.47$; $p < .05$), felt it was more successful ($F=6.04$; $p < .05$) and liked REA more ($F=3.24$; $p < .08$), and that REA was more intelligent ($F=3.40$; $p < .08$), and knew them better ($F=3.42$; $p < .08$) than active subjects. These differences may be explained by the fixed-initiative dialogue model used in the wizard-of-oz script. REA's interaction was designed for passive users--there was very little capability in the interaction script to respond to unanticipated user questions or statements--and user initiation attempts were typically met with uncooperative system responses or interruptions. But, given the choice between phone and face-to-face, passive users preferred to interact with REA face-to-face: they rated her as more friendly ($F=3.56$; $p < .07$) and informed ($F=6.30$; $p < .05$) in this condition. Passive users also found the phone to be more tedious, while active users also found the phone to be less tedious ($F=5.15$; $p < .05$). Active users may have found the face-to-face condition particularly frustrating since

processing delays may have led to the perception that the floor was open (inviting an initiation attempt), when in fact the wizard had already instructed REA to produce her next utterance.

However, when interacting on the phone, active users differed from passive users in that active users felt she was more reliable when using social dialogue and passive users felt she was more reliable when using task-only dialogue. When interacting face-to-face with REA, there was no such distinction between active and passive users ($F=4.67$; $p<.05$).

5.3.2 Effects on Task Measure

One of the most interesting results obtained is that extroverts were willing to pay more for the same apartment in the embodied condition, while introverts were willing to pay more over the phone ($F=3.41$; $p<.08$), as shown in Figure 5-5.

While potentially very significant, this finding is a little difficult to explain, especially given that trust did not seem to play a role in the evaluation. Perhaps, since subjects were asked to simply play the role of someone looking for an apartment, and given that the apartments displayed were cartoon renditions, the subjects may not have felt personally invested in the outcome, and thus may have been more

likely to be persuaded by associative factors like the perceived liking and credibility of REA. In fact, trust has been shown to not play a role in persuasion when "peripheral route" decisions are made, which is the case when the outcome is not of personal significance (Petty & Wegener, 1998). Further, extroverts are not only more sociable, but more impulsive than introverts (Wilson, 1977), and impulse buying is governed primarily by novelty (Onkvisit & Shaw, 1994). Extroverts did rate face-to-face interaction as more engaging than phone-based interaction (though not at a level of statistical significance), while introverts rated phone-based interactions as more engaging, providing some support for this explanation. It is also possible that this measure tells us more about subjects' assessment of the house than of the realtor.

5.3.3 Gender Effects

Women felt that REA was more efficient ($F=5.61$; $p<.05$) and reliable ($F=4.99$; $p<.05$) in the embodied condition than when interacting with her over the phone, while men felt that she was more efficient and reliable by phone. Of course, REA has a female body and a female voice and so in order to have a clearer picture of the meaning of these results, a similar study would need to be carried out with a male realtor.

5.3.4 Qualitative Results

Following the interaction with REA and filling out their self-report questionnaires, subjects were briefly interviewed about their experience.

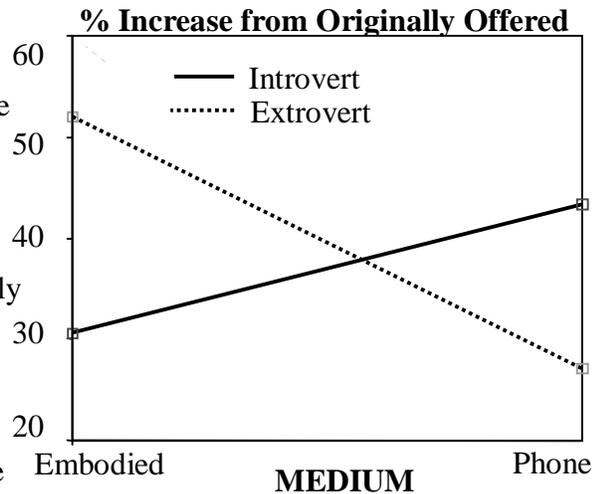


Figure 5-5. Amount Subjects Were Willing To Pay

Overall Impressions

Overall, subjects seemed very comfortable interacting with REA and were able to readily engage in conversation with her, even though many reported that it seemed a little strange at first:

I thought it was very cool.

I felt comfortable enough. I just felt it was the wave of the future.

Getting used to her voice was a little awkward at first, but other than that no, it was pretty natural.

Several subjects commented that REA's nonverbal behavior seemed a little strange, especially her gaze-away behavior between turns:

...when she looked to the side and then back before saying something was little bit unnatural.

But, as a person, her looking away, that wasn't good.

When she paused before responding, or looked away before bringing up a new topic, that seemed more like a computer.

Several subjects also commented on REA's lack of uptake on their conversational moves, or apparent lack of understanding by showing them apartments that did not meet their criteria (which may have had a negative impact on trust):

The thing that really stood out was that I felt like she didn't really [understand], I said something about walkspace or warehouse space studio, and she didn't come up with anything that really matched that.

Some of the time I don't think REA understood what I was asking.

And I think I asked it "Does that look over the Charles?" and it just said "yea." "DO YOU UNDERSTAND ME?" it said yes, and I saw from the graphic that it looked out over Boston.

Small Talk and Trust

Opinions on the naturalness and liking of REA's small talk (from those in the SOCIAL condition) ranged widely. Several subjects reported liking the social dialogue aspects of the interaction:

It was amazingly normal, given that it was a computer. It was really well done. I was thinking that if she can do it, then any person can learn how to chit chat.

It wasn't just real estate talk, so I felt like it made her more human.

A lot of her small talk was joking about how ridiculous the rents are in Boston, which is sort of something as a buyer, you kinda like to hear. It sounds like she's on your side when she says things are expensive. It didn't come through as stilted, it was kinda funny, and it didn't get in the way of things, because it seemed that it was between looking at a place and looking at the one and seemed perfectly natural.

It was more engaging than talking with a human. I was chatting near the end.

One subject in SOCIAL condition even thought that REA did not do enough small talk:

She maybe still seems a little bit cold. For something like renting an apartment I guess I'm looking for a lot more up front of who are you, what are you looking for. She asked a lot of what kind of style are you looking for, very much like let's get down to business, let me show you some places. It's like she's too busy to find out about me, she just wants to get out there and show me some places.

And, there were comments from several subjects in TASK condition who felt that REA did not do enough rapport building:

Yes, she asked the right questions about what I wanted, but in terms of getting to know her, that didn't take place. I felt like I was talking to a machine vs. talking to a person. That was the only thing that kind of threw me off. I would have liked her to ask more questions about what I like, getting to know more who I am, that would have made me more comfortable, at least in this scenario. Buying groceries would be a different thing. But for real estate, it's a very personal thing, at least for me, to feel comfortable with who I'm buying it from, who's selling it to me, and the place that I'm buying. She didn't interact with me enough as an individual.

Other subjects in SOCIAL condition clearly did not like REA's small talk at all, although it was often unclear if it was because they did not like small talk in general or because it was a computer doing it:

I really got a queasy feeling from her chit chat, because it just seems so insincere to have a machine acting like your buddy. It was creepy.

But initially, when she spent a minute or two minutes, with all the BS about the weather and shit, that was like a really unpleasant reminder that this wasn't a person. It just felt so manipulative, the whole "How you doin'?" That's what I felt.

Her conversational style, seemed to me to be more applicable to women, frankly, than to me. I come in and I shop and I get the hell out. She seemed to want to start a basis for understanding each other, and I would glean that in terms of our business interaction as compared to chit chat.

As one subject observed, having small talk with a computer involves a certain level of pretense, but then small talk itself is typically full of pretense:

I guess it was a little contrived, but pretty normal in terms of the kind of small talk. But I guess small talk is contrived.

Is REA a Person?

Finally, subjects were asked at the end of the experiment if they thought they were interacting with a person or a computer (the question was initially intended as a manipulation test to see if they suspected the wizard-of-oz deception). Surprisingly, subjects classified REA closer to a person than a computer, but with some hedging and qualification:

More of a person, especially with the social dialog.

It felt like a computer with personality. Sometimes it felt like a computer, like when it didn't understand me, but other times when the conversation was going well it felt like a person.

A computer-person I guess. It was a lot like a human.

5.4 Discussion

This study provided some important results and lessons learned to be incorporated into the development of future relational agents. Most importantly, it demonstrated that people

will readily engage in relational dialogue (or at least small talk) with an ECA; they do not find it uncomfortable or disorienting, and they will do so with no prior training. While many did not like conducting small talk, many did, and found this to be an important part of their expectations for dealing with a real estate agent. There is also some evidence that small talk can lead to increased trust, although in this study this occurred only for a subset of subjects (extroverts) and in one medium (face-to-face with an ECA).

The negative outcomes of this study, while more numerous than the positive ones, provide even more important lessons for what does, and does not, work for a relational agent.

First, REA's nonverbal behavior was entirely inappropriate for social dialogue, as reflected in subjects preferring to conduct small talk with her on the phone, while still preferring task-only talk face-to-face. None of the nonverbal behaviors associated with warmth and immediacy were implemented: REA had no functional facial expressions beyond a gaze-away (which itself was flawed) and occasional smile; she did not mark social dialogue with increases in gaze, gestural and facial animation and decreased proximity. These problems are partially addressed by the system discussed in the following chapter.

Second, REA's overall visual persona was very cold and uninviting, as reflected in subjects giving her higher ratings on friendliness, warmth and relational closeness on the phone, compared with face-to-face. Her lack of facial expressions, frequent sustained gaze-aways, stiff posture, and overall robotic animation also led to what appeared to be an assessment of introversion by subjects. These can be difficult hurdles for a relational agent to overcome in such a brief interaction.

Third, no amount of relational behavior can compensate for incompetence. REA's fixed script and limited stock of two apartments led many subjects to make negative comments about her during the debrief interview when either she could not answer their questions or could not show them apartments that met their criteria. The application logic underlying a relational agent must be correct, otherwise social dialogue can come off as insincere and duplicitous. Given the errors that will inevitably arise in natural language interfaces, an agent can attempt to minimize these problems by clearly stating what its capabilities and limitations are at the outset.

Fourth, REA's lack of uptake of subjects' social dialogue moves and inability to answer their questions led many subjects to believe that, while she was correctly executing the ritual of small talk (e.g., with backchannel responses such as "Really? Cool."), at a deeper level she really wasn't understanding what they were saying (based on the debrief comments in the previous section). The combination of a speech interface, conversational setting, and social dialogue yields a context in which a person feels entirely unconstrained in what they can say and how they say it, making the development of dialogue systems for these contexts AI- and NLP-complete. Until these problems are solved, social dialogue will only work in applications in which the setting and dialogue context are crafted to greatly limit what can naturally be said without destroying the relaxed feel of social dialogue (a very challenging problem), cues to what can be said at any given time are provided to the user, or speech input is replaced with dynamic menus in which the allowable inputs are clearly enumerated.

Fifth, real-estate sales, while providing an application context in which social dialogue is acceptable, is a poor application area for evaluating the real impact relational agents can have on people's lives. It will be difficult to find subjects who are actually willing to purchase a house from an ECA, let alone do it in a laboratory setting for the benefit of science. Until such

subjects are found, people must be asked to play “what if” games in these scenarios, and their behavior may be completely unrelated to what they would do in real life (as one subject put it “It was kinda fun to fantasize about having \$3000, an apartment of my choice”). In addition, many people are predisposed to not trust real estate agents, and this provides a situation in which building trust and rapport is an uphill battle (“There’s two things that come to mind when I think of somebody I can’t trust—and that’s a real estate agent and a lawyer.”). However, real-estate sales may still be a good domain in which an ECA can give users information about properties—just as people use web sites to shop for homes—as long as the task does not involve asking people to actually commit to a purchase.

Finally, relative to the theory presented in chapter 3, it should be clear that very little can be accomplished relationally in a single five minute conversation, beyond initial attraction and first impression effects. In such a brief scenario, the best that a relational agent can hope to do is to establish expectations of the stereotypical relationship that will be drawn from most heavily in its interactions with the user, and to uphold those expectations through appropriate relational dialogue.

5.5 Summary

This chapter presented an evaluation of the use of social dialogue by an embodied conversational agent and its effect on subjects’ trust in the agent in a single face-to-face conversation. Some evidence was found for increasing trust for one group of users (extroverts) in one medium (face-to-face), but this result may have been due to other effects. A number of important lessons were also learned for the development of future relational agents, such as the need to use nonverbal behavior appropriate to the type of interaction (e.g., social vs. task). The following chapters outline the development of a relational agent that addresses many of these shortcomings.

Chapter 6

GENERATING NONVERBAL BEHAVIOR FOR FACE-TO-FACE RELATIONAL DIALOGUE

Most relational communication is conducted in face-to-face conversation, where the myriad social cues afforded by the face, hands and body provide important information about interlocutors' attitudes towards each other, in addition to providing additional channels of communication for propositional and affective information, and cues to help regulate the structure of the interaction. This chapter presents a computational model for automatically generating many of these nonverbal cues. The model is an extension of the BEAT text-to-embodied speech system, developed in collaboration with Hannes Vilhjálmsón, and reported in (Cassell, Vilhjálmsón, & Bickmore, 2001).

6.1 Introduction

Why is it that people build and maintain relationships primarily in the context of face-to-face interaction? Several studies have demonstrated what most business people already know; when the social aspects of an interaction are especially important--such as when you are getting to know a new client or negotiating an important deal--nothing beats face-to-face interaction. In a review of studies comparing video and audio-mediated communication, Whittaker and O'Conaill concluded that video was superior to audio-only for social tasks while there was little difference in subjective ratings or task outcomes in tasks in which the social aspects were less important (Whittaker & O'Conaill, 1997). They found that for social tasks, interactions were more personalized, less argumentative and more polite when conducted via video-mediated communication, that participants believed video-mediated (and face-to-face) communication was superior, and that groups conversing using video-mediated communication tended to like each other more, compared to audio-only interactions.

Obviously, some nonverbal communication must be responsible for these differences. In this chapter I look at the role of nonverbal behavior in building and maintaining relationships in the context of face-to-face conversation, and present a computational model that provides some of this functionality for embodied conversational agents.

6.2 The Role of Nonverbal Behavior in Face-to-Face Interaction

According to Argyle, nonverbal behavior is used to express emotions, to communicate interpersonal attitudes, to accompany and support speech, for self presentation, and to engage in rituals such as greetings (Argyle, 1988). Of these, coverbal and emotional display behaviors have received the most attention in the literature on embodied conversational agents and facial and character animation in general (e.g., (Cassell, Sullivan, Prevost, & Churchill, 2000)). Coverbal behavior can be further broken down into behaviors that convey information (support the "propositional" functions of face-to-face conversation, such as pointing at an object being discussed) and those which regulate aspects of the interaction (support the

"interactional" or "envelope" functions of face-to-face conversation, such as gazing away from one's interlocutor when beginning a speaking turn).

A given communicative behavior, such as raising one's eyebrows, can be used in a variety of contexts to achieve different effects, such as signaling surprise or emphasis. In addition, a given communicative function, such as emphasis, can be achieved by a variety of behaviors, such as raising one's eyebrows or gesturing. An agent trying to achieve a particular communicative goal must know the range of behaviors available and be able to substitute one for another when necessary (e.g., emphasizing with his eyebrows if his hands are occupied). Further, the mapping from communicative function to behavior is culture-dependent. Thus, in developing models of nonverbal behavior encoding and decoding, it is important to separate communicative functions from behaviors.

In the following sections I will discuss the major functions of nonverbal behavior in face-to-face interaction, with particular emphasis on coverbal behaviors, splitting this latter category into interactional and propositional functions. The notion of conversational frames is also introduced, and since these can serve both coverbal propositional and interactional functions they are described separately.

6.2.1. Emotional Display Functions

As discussed in Chapter 3, emotion plays a crucial role in relationships. In face-to-face relational conversation, emotional display plays a particularly important role in communicating emotional state to one's partner. Although there have been several hypothesized mechanisms by which emotions can be generated (Izard, 1993), it is clear that emotional display is the result of both spontaneous readouts of internal state (a "conversation between limbic systems") and deliberate communicative action (Buck, 1991), and many researchers even believe the primary function of emotions is the communication of feelings and needs to others (Andersen & Guerrero, 1998). Empathy and emotional support processes--crucial in all close relationships--cannot take place without an antecedent emotional state being communicated.

Emotional display can be a very complex phenomenon involving a wide range of verbal and nonverbal behavior. Planalp performed an experiment in which she asked subjects to monitor a person they knew well, and when they felt that this person was having an emotion they were to document what emotion it was and how they could tell (Planalp, 1998). Almost all subjects used multiple cues (mode=4), including (in decreasing order of frequency)¹:

- Vocal cues (loudness, speed of talking, amount of talking)
- Indirect verbal cues (name-calling, apologizing, long discourses describing the emotion)
- Facial cues (eye gaze, smiling, movements of lips or eyebrows, rolled eyes, puffy eyes)
- Activity cues (solitary - going for a walk, eating, drinking, smoking, throwing things; interpersonal - kissing, hugging, tickling, hitting, touching)
- Context cues (knowing the person had just received a letter or phone call, was getting married, had a school assignment due, was attending a hearing, etc.)

¹ The terminology used here is Planalp's.

- Body cues (clenched fists, walking heavily, stamping feet, throwing arms up, stomping around)
- Trait cues (reference to typical behaviors: the person typically shows this expression when emotional; or is behaving atypically)
- Physiological cues (crying, blushing, changes in breathing)
- Direct verbal ("I'm happy", "He pissed me off")

In addition to the surprising range of cues, analysis of her examples shows that there may not be any simple calculus for combining these cues; they unfold in complex ways over a period of time.

Even though emotions are often decoded using such a wide range of cues, there does appear to be some primacy to emotional displays involving facial expression. Ekman conducted a study of facial displays across a number of cultures and found that there is strong evidence for a "basic" set of emotions including fear, anger, disgust, sadness and joy, since these seem to occur and have the same meaning in a wide range of cultures (Ekman, 1993). In addition, communicative emotional display in face-to-face conversation may also rely primarily on the facial expression because the face is so prevalent in this context and is being continuously monitored for other propositional and interactional communicative cues.

Thus, relational agents should have the ability to generate (and ultimately recognize) emotional display in face-to-face conversations with users. The range of displays depends on the application area, but should, at a minimum, include displays of concern and positive and negative valence (happy/sad) to accompany empathetic dialogue, if the agent is to engage in empathetic behavior (positively and negatively valenced displays are the minimum required to demonstrate understanding of emotionally-valenced messages, and a concerned display indicates the activity that the listener is engaged in).

6.2.2 Attitudinal Functions

In addition to displaying affect, and co-verbal and ritual behavior, the most important use of nonverbal behavior in relationship building and maintenance is the display of interpersonal attitude (Argyle, 1988). The display of positive or negative attitude can greatly influence initial perceptions of people we meet and whether we approach them or not.

The most consistent finding in this area is that the use of nonverbal "immediacy behaviors" (also called affiliative or liking behaviors)--close conversational distance, direct body and facial orientation, forward lean, increased and direct gaze, smiling, pleasant facial expressions and facial animation in general, nodding, frequent gesturing and postural openness--projects liking for the other and engagement in the interaction, and is correlated with increased solidarity (Argyle, 1988; Richmond & McCroskey, 1995).

In one of the most famous and mis-quoted studies on the nonverbal encoding of attitude, Mehrabian conducted a series of studies in which he attempted to determine the relative effects of facial expression, intonation and speech on a hearer's perception of a speaker's attitude towards them (Mehrabian, 1972). In this study recordings were made of a single positive, neutral or negative word, read by three women with an imagined liking, neutral or disliking attitude towards their listener. Photos of three female models were also taken as they portrayed like, neutrality or dislike of their imagined addressee. Subjects then heard one recording and saw one picture and rated the attitude of the speaker towards them. Mehrabian then did a regression to determine contribution of each channel to the subjects' rating of the

speakers' like or dislike of them. The results indicated that 55% of the attitudinal message was conveyed via facial expression, 38% via intonation, and only 7% via verbal content. His conclusion was that if there is a conflict between attitudinal cues, people tend to derive more information from visual or tonal cues in order to disambiguate meaning.

Other studies of attitudinal communication indicate that people are very good at decoding attitudinal messages. They also do not like inconsistent verbal and nonverbal messages, and find such speakers insincere. Finally, when any channel is conveying negative attitude, people tend to assess the overall message as negative in tone (Argyle, 1988).

6.2.3 Coverbal Interactional Functions

Table 6-1 shows examples of mappings from interactional communicative functions to behaviors and is based on previous research on typical North American nonverbal displays (Cassell, Nakano, Bickmore, Sidner, & Rich, 2001; Duncan, 1974; Goodwin, 1981).

Communicative Functions	Communicative Behavior
<i>Turn-Taking:</i>	
Give Turn	Paralinguistic drawl on final syllable of clause Termination of hand gesture Completion of clause
Keep Turn	Gaze away
Take Turn	Gaze away Start of hand gesture
<i>Feedback:</i>	
Request Feedback	Gaze towards & End clause Pause or Restart
Give Feedback	Looking, Head Nod
<i>Discourse segment changes:</i>	
Topic shift	Posture shift

Table 6-1. Some examples of conversational functions and their associated behaviors

Conversational Turn-Taking and Interruption

Interlocutors do not normally talk at the same time, thus imposing a turn-taking sequence on the conversation. The protocols involved in floor management -- determining whose turn it is and when the turn should be given to the listener -- involve many factors including gaze and intonation. In addition, listeners can interrupt a speaker not only with voice, but by gesturing to indicate that they want the turn.

Feedback and Error Correction

During conversation, speakers can non-verbally request feedback from listeners through gaze and raised eyebrows and listeners can provide feedback through head nods and

paraverbals ("uh-huh", "mmm", etc.) if the speaker is understood, or a confused facial expression or lack of positive feedback if not. The listener can also ask clarifying questions if they did not hear or understand something the speaker said.

Topic Shifts

Changes in discourse structure can be indicated by discourse markers (e.g., "anyway", "first", etc., (Schiffrin, 1987)), but also nonverbally by posture shifts--large movements in limbs or stance (Cassell, Nakano et al., 2001).

6.2.4 Coverbal Propositional Functions

Hand gestures can convey propositional information that is either redundant with, or complementary to, the propositional content of the accompanying speech (McNeill, 1992). Probably the most commonly thought of use of the body in conversation is the pointing (deictic) gesture, possibly accounting for the fact that it is also the most commonly implemented use for the bodies of animated interface agents. In fact, however, most conversations don't involve many deictic gestures unless the interlocutors are discussing a shared task that is currently present. Typically, the most common conversational gesture is the "beat" or baton gesture, which is a formless motion of the hand synchronized with the part of speech that is being emphasized. Iconic and metaphoric gestures convey some features of the action or event being described. Whereas iconics convey information about spatial relationships or concepts, metaphorics represent concepts which have no physical form, such as a sweeping gesture accompanying "the property title is free and clear."

Other examples of nonverbal behaviors that can convey propositional information are eyebrow raises (when used for emphasis) and, of course, emotional displays.

6.2.5 Coverbal Framing Functions

Frame changes are similar to topic shifts (described above) but carry even more information about the type of interaction that is being initiated. Tannen provides an excellent summary of the history and varying conceptualizations of the term "frame" as used in sociolinguistics (Tannen, 1993a, 1993b). In her review, Tannen groups this work into two broad types: *interaction frames*, which represent what people think they are doing when they talk to each other (e.g., small talk vs. negotiation vs. job interview); and *knowledge schema* which are expectations associated with situations, objects, people, etc. Gregory Bateson introduced the notion of *frame* in 1954, and showed that no communication could be interpreted without a metamessage about what was going on, i.e., what the frame of interaction was. He showed that even monkeys exchange signals that allow them to specify when the "play" frame is active so that hostile moves are interpreted in a non-standard way (Bateson, 1954). Charles Fillmore defined frame as any system of linguistic choices associated with a scene (where a scene is any kind of coherent segment of human actions) (Fillmore, 1976). Gumperz described this phenomena (he called contextualization) as exchanges representative of socioculturally familiar activities, and coined "contextualization cue" as any aspect of the surface form of utterances which can be shown to be functional in the signaling of interpretative frames (Gumperz, 1977).

Most conversational systems developed to date operate in a single, task-oriented interactional frame, characterized by maximum Gricean cooperativeness, and thus do not need

to represent multiple frames nor worry about how frame changes are signaled. However, relational agents need to be able to conduct, at a minimum, relational dialogue (such as small talk) and task-oriented talk, and thus require the ability to clearly signal to the user when a change between these frames has occurred. The reason this is important for a model of relational nonverbal behavior is that many, if not most, contextualization cues are nonverbal in nature. Gumperz describes a number of examples of frame changes signaled by prosody (intonation and stress) and paralinguistics (e.g., pitch register, rhythm, loudness, etc.) and goes on to define contextualization cue as "any aspect of the surface form of utterances which, when mapped onto message content, can be shown to be functional in the signaling of interpretative frames" (Gumperz, 1977). Goodwin provides several more examples of contextualization cues, including emotional displays, smiling, laughing and posture shifts (Goodwin, 1996).

In the implementation that follows, I implement four conversational frames—TASK, SOCIAL, EMPATHY and ENCOURAGE—and use facial expression, proximity and speech rate as contextualization cues. While frame can be nested in natural interaction (e.g., SOCIAL conversation within a TASK interaction) I currently only take into account a single frame's influence on nonverbal behavior.

6.3 BEAT

BEAT is a text to embodied speech translation system that I co-developed with Hannes Vilhjálmsón at the Media Lab, based on an idea by Justine Cassell² (Cassell, Vilhjálmsón et al., 2001). BEAT takes the text of an utterance as input (optionally tagged with semantic and pragmatic markers) and produces an animation script as output that can be used to drive an ECA's production of the utterance, including not only speech and intonation, but accompanying nonverbal behavior, such as hand gestures, gaze behavior, and eyebrow raises. BEAT was initially developed with a small set of interactional and propositional nonverbal behaviors that it could generate, but was developed to be extensible so that new conversational functions and behaviors could be easily added. This section describes the basic BEAT system and its propositional and interactional coverbal behaviors; section 6.4 then presents extensions that were made to BEAT to support relational dialogue.

The BEAT system is built to be modular and to operate in real-time and, as mentioned, to be easily extensible. To this end, it is written in Java, is based on an input-to-output pipeline approach with support for user-defined filters and knowledge bases, and uses XML as its primary data structure. Processing is decomposed into modules that operate as XML transducers; each taking an XML object tree as input and producing a modified XML tree as output. The first module in the pipeline operates by reading in XML-tagged text representing the text of the character's script and converting it into a parse tree. The various knowledge

² Hannes and I collaborated on the overall BEAT architecture and the behavior suggestion and selection framework. He designed and implemented the knowledge base, language tagging module and discourse model. I designed and implemented the XML library, the behavior generators and filters described in 6.3.4 and 6.3.5, the behavior scheduling module, and translators for the Pantomime animation system, animator dope sheets, McNeill format display, and Gantt chart display.

bases used in the system are also encoded in XML so that they can be easily extended for new applications.

New pipeline XML transducers, as well as nonverbal behavior generators and filters (discussed in sections 6.3.3 and 6.3.4) can be authored through Java subclassing to facilitate extensibility. The system is real-time in that the time to produce an utterance is typically less than the natural pause between speaker turns in a dialogue (typically between 500 -1000ms). This is enabled by the pipeline architecture in which all operations are performed on a single XML tree within a single Java program.

XML provides a natural way to represent information which spans intervals of text, and its use facilitates modularity and extensibility by allowing users to add their own tags to the parse tree at any stage of processing. The combination of XML and Java also provide cross-platform portability, since both have been designed with this as a primary design goal. Nonverbal behavior generators and filters can also be authored in XSL, an XML-based scripting language, which provides extensibility without having to program in Java. The use of a validating XML parser enables automatic testing of the output from each module during development. There are also many tools available for parsing, generating and displaying XML, which provide great leverage during system development.

An overview of the system is shown in Figure 6-1. There are three main processing modules: Language Tagging, Behavior Generation and Behavior scheduling. The stages of XML translation produced by each of these modules are shown in Figure 6-2. The Behavior Generation module is further divided into a Suggestion module and a Selection module, as the approach to the generation process is to first suggest all plausible behaviors and then use user modifiable filters to trim them down to a set appropriate for a particular agent. In Figure 6-1, user definable data structures are indicated with dotted line boxes. I will now discuss each of these components in turn.

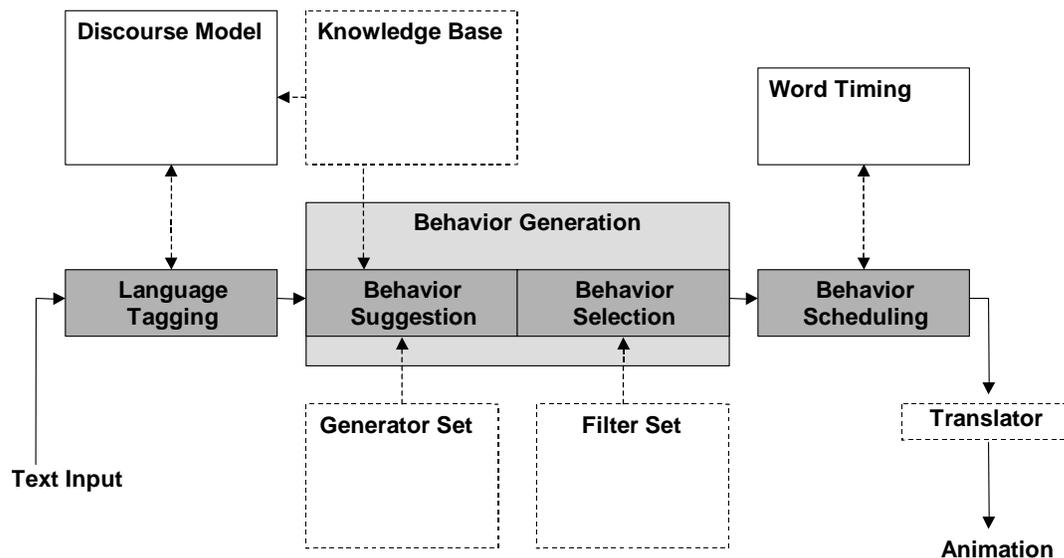


Figure 6-1. BEAT System Architecture

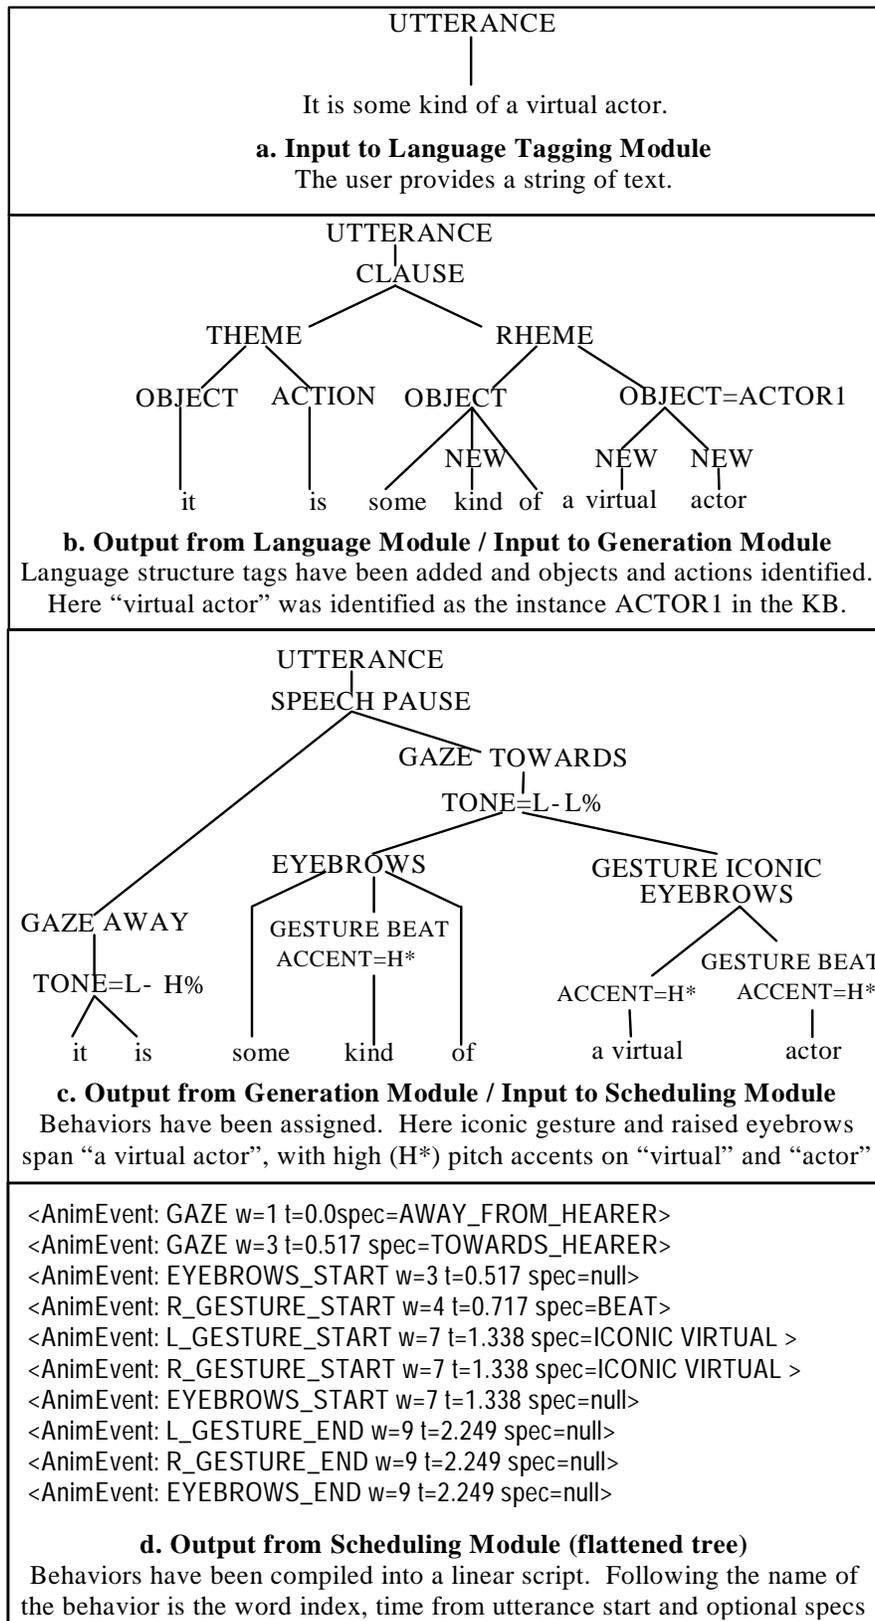


Figure 6-2. XML Trees Passed Among Modules

6.3.1 Knowledge Base

A knowledge base adds some basic knowledge about the world to what can be understood from the text itself, and therefore allows inferences to be drawn from the typed text, and consequently the kinds of gestures that should illustrate it, and the kinds of places where emphasis should be created. Currently, the knowledge base is stored in two XML files, one describing objects and other describing actions. These knowledge bases are seeded with descriptions of generic objects and actions but can easily be extended for particular domains to increase the efficacy of nonverbal behavior assignment.

The object knowledge base contains definitions of object types and instances of those types. Figure 6-3 shows three example entries. The first defines a new object type PROFESSIONAL as of the person class (vs. object or place) with symbolic features such as TYPE, describing whether the professional is REAL or VIRTUAL; and ROLE, describing the actual profession. For each feature, typical values are described (e.g., real professionals are typical, while virtual ones are not), which is important since people tend to generate iconic gestures for the unusual aspects of objects they describe (Yan, 2000). The second knowledge base entry defines an object instance and provides values for each feature defined for the type. The last entry is a description of a gesture that could represent the value VIRTUAL.

The action knowledge base contains associations between domain actions and hand gestures that can depict them. An example entry is:

```
<GESTURE NAME="MOVE" TYPE="ICONIC">  
<RIGHTARM HANDSHAPE=5, TRAJECTORY= "moves from CC towards L ..."/>  
</GESTURE>
```

which simply associates a particular gesture specification with the verb to move.

<pre><TYPE NAME="PROFESSIONAL" CLASS="PERSON"> <SYMFEATURE NAME="ROLE" TYPICAL="ANY"/> <SYMFEATURE NAME="TYPE" TYPICAL="REAL"/> <NUMFEATURE NAME="AGE" TYPICAL="25-65"/> </TYPE></pre>
<pre><INSTANCE OF="PROFESSIONAL" ID="ACTOR1" ROLE="ACTOR" TYPE="VIRTUAL" "/></pre>
<pre><GESTURE TYPE="ICONIC" VALUE="VIRTUAL"> <RIGHTARM HANDSHAPE="virtual" TRAJECTORY="virtual"/> </GESTURE></pre>

Figure 6-3. Example KB entries that describe an instance of a professional, that surprisingly is virtual – an attribute that has a defined gesture form.

As mentioned above, the system comes loaded with a generic knowledge base, containing information about some objects and actions, and some common kinds of gestures with prototypical form. Those common gestures include the beat, which is a formless flick of the hand, the deictic, which is a pointing gesture, and the contrast gesture (see Section 6.3.4). The other major kind of gesture, an iconic, represents some object or action, and may be performed differently by different speakers and in different contexts. These gestures are added to the database by the animator. All gestures are specified using a compositional notation in which hand shapes and arm trajectories for each arm are specified independently. This makes the addition of new gestures easier, since existing trajectories or hand shapes can be re-used.

6.3.2 Discourse Model

The BEAT discourse model currently keeps track of the root forms of words spoken by the agent in past utterances, so that NEW words can be detected and tagged as well as words that are known to contrast with previously spoken words (using WordNet). The location of NEW words in an utterance is used by the language tagging module to determine the breakdown of clauses into THEME and RHEME, as well as by several of the behavior generators described in section 6.3.4. Ultimately, the discourse model should keep track of previously mentioned entities for proper determination of THEME and RHEME (the current approach uses a set of heuristics, described in the following section).

6.3.3 Language Tagging

The language module of BEAT is responsible for annotating input text with the linguistic and contextual information that allows successful nonverbal behavior assignment and scheduling. It should be noted that much of what is described in this section is similar to or, in some places identical, to the kind of tagging that allows TTS systems to produce appropriate intonational contours and phrasing along with typed text (Hirschberg, 1990). Additional annotations are used here, however, to allow not just intonation but also facial display and hand gestures to be generated. And, these annotations will allow not just generation, but also synchronization and scheduling of multiple nonverbal communicative behaviors with speech.

The largest unit is the UTTERANCE, which is operationalized as an entire paragraph of input. The utterance is broken up into CLAUSES, each of which is held to represent a proposition. To detect clause boundaries the tagging module looks for punctuation and the placement of verb phrases.

Clauses are further divided into two smaller units of information structure, a THEME and a RHEME. The former represents the part of the clause that creates a coherent link with a preceding clause and the latter is the part that contributes some new information to the discussion (Halliday, 1973). For example in the mini-dialogue "who is he?" "he is a student", the "he is" part of the second clause is that clause's theme and "student" is the rheme. Identifying the rheme is especially important in the current context since gestural activity is usually found within the rheme of an utterance. The language module uses the location of verb phrases within a clause and information about which words have been seen before in previous clauses to assign information structure, following the heuristics described in (Hiyakumoto, Prevost, & Cassell, 1997).

The next to smallest unit is the word phrase, which in the current implementation either describes an ACTION or an OBJECT. These two correspond to the grammatical verb phrase and noun phrase, respectively. Actions and objects are linked to entries in the knowledge base whenever possible, as follows. For actions, the language module uses the verb head of the corresponding verb phrase as the key to look up an action description in the action database. If an exact match for that verb is not found, it is sent to an embedded word ontology module (using WordNet (Miller, Beckwith, Fellbaum, Gross, & Miller, 1993)), which creates a set of hypernyms and those are again used to find matching descriptions in the knowledge base. A hypernym of a word is a related, but a more generic -- or broader -- term. In the case of verbs, one can say that a certain verb is a specific way of accomplishing the hypernym of that verb. For example “walking” is a way of “moving”, so the latter is a hypernym of the former. Expanding the search for an action in the action database using hypernyms makes it possible to find and use any descriptions that may be available for a super-class of that action. The database therefore doesn’t have to describe all possible actions, but can focus on high-level action categories. When an action description match is found, a description identifier is added to the ACTION tag.

For objects, the module uses the noun head as well as any accompanying adjectives to find a unique instance of that object in the object database. If it finds a matching instance, it adds the unique identifier of that instance to the OBJECT tag.

The smallest units that the language module handles are the words themselves. The tagger uses the EngLite parser from Conexor (www.conexor.fi) to supply word categories and lemmas for each word. The module also keeps track of all previously mentioned words and marks each incoming noun, verb, adverb or adjective as NEW if it has not been seen before. This “word newness” helps to determine which words should be emphasized by the addition of intonation, eyebrow motion or hand gesture (Hiyakumoto et al., 1997).

Words can also stand in contrast to other words (for example “I went to buy red apples but all they had were green ones”), a property often marked with hand gesture and intonation and therefore important to label. The language module currently labels contrasting adjectives by using WordNet to supply information about which words might be synonyms and which might be antonyms to one another (Hiyakumoto et al., 1997). Each word in a contrast pair is tagged with the CONTRAST tag.

In sum, the language tags that are currently implemented are:

- Clause
- Theme and rheme
- Word newness
- Contrast
- Objects and actions

6.3.4 Behavior Suggestion

The Behavior Suggestion module operates on the XML trees produced by the Language Tagging module (such as the one shown in Figure 6-2b) by augmenting them with suggestions for appropriate nonverbal behavior. This augmentation is intended to be liberal and all-inclusive; any nonverbal behavior that is possibly appropriate is suggested independent of any other. The resulting over-generated behaviors will be filtered down in the next stage of processing to the final set to be animated. This independence of behavior

suggestions allows filters to be defined for different personality types, situations, and scenes (for example, an animator may choose to filter out fewer gestures when animating the effusive bubbly personality than when animating the taciturn introvert).

Behavior suggestion proceeds by applying each of an extensible set of nonverbal behavior generators to all nodes in the XML tree which meet criteria specified by each generator. When the criteria are completely satisfied a suggestion is added to the appropriate node. The pseudocode for the generator which suggests beat gestures is shown in Figure 6-4 (behavior generators are actually implemented in Java).

```
FOR each RHEME node in the tree
  IF the RHEME node contains at least one NEW node
  THEN Suggest a BEAT to coincide
       with the OBJECT phrase
```

Figure 6-4. Example Behavior Generator

This pseudocode states that beat gestures are appropriate during the description of objects (noun phrases), but only when those objects are part of the rheme (new information) and contain new words.

Behavior suggestions are specified with a tree node (defining the time interval they are active for), priority (used for conflict resolution), required animation degrees-of-freedom, and any specific information needed to render them (e.g., gesture specification). Suggestions also specify whether they can co-articulate, i.e., occur during other behaviors which use the same degrees of freedom. For example, beat gestures can co-articulate with other gestures through the addition of a relative hand displacement (Cassell et al., 1994).

The current set of behavior generators implemented in the toolkit includes the following:

Beat GestureGenerator

Beats, or formless handwaves, are a "default" gesture, in that they are used when no additional form information is available to generate a more specific kind of gesture, and they account for roughly 50% of the naturally occurring gestures observed in most contexts (McNeill, 1992). Thus, they are typically redundantly generated when other types of gestures are appropriate, but they are given a low priority relative to other types of gestures so that they will only be selected when no other gestures are available. Like all gestures that occur during speech, beats occur primarily during the introduction of new material (rheme).

Surprising Feature Iconic Gesture Generator

A study of individuals describing house floor plans showed that gestures representing some feature not described in accompanying speech were used 80% of the time during the description of house features which were "surprising" or unusual in some way, (Yan, 1999). Following these results, this generator determines if any of the OBJECTS identified by the Tagger within the RHEME have unusual features (based on information in the object knowledge base), and for each generates an iconic (representational) gesture based on the gesture specification defined on the unusual feature value in the knowledge base.

Action Iconic Gesture Generator

This generator determines if there are any actions (verb phrase roots) occurring within the RHEME for which gestural descriptions are available in the action knowledge base. For each such action, an iconic gesture is suggested with the gesture specification used from the knowledge base.

Contrast Gesture Generator

The tagger identifies objects which contrast with other nearby objects (e.g., "I don't know if this is a good thing or a bad thing."). Such objects (even if they occur within a THEME) are typically marked with either beats or a "contrastive gesture" if there are exactly two such objects being contrasted (gestures literally of the form "on the one hand...on the other hand") (Cassell & Prevost, 1996). This generator suggests beats for contrast items unless there are exactly two items being contrasted, in which case the special contrast gesture is suggested.

Eyebrow Flash Generator

Raising of eyebrows can also be used to signal the introduction of new material (Pelachaud, Badler, & Steedman, 1994). This generator suggests raising the character's eyebrows during the description of OBJECTs within the RHEME.

Gaze Generator

(Torres, Cassell, & Prevost, 1997) studied the relationship between eye gaze, theme/rheme, and turn-taking, and used these results to define an algorithm for controlling the gaze behavior of a conversational character. The gaze generator that implements this algorithm is shown in Fig. 6-5. (Relative to the problematic gaze behavior used in the REA small talk experiment—described in chapter 5—this algorithm allows for changes in gaze direction during a turn, yielding a much more natural interaction. The model used in the REA wizard-of-oz control used a simple gaze-away, pause, gaze-towards at the start of every turn.)

```
FOR each THEME
    IF at beginning of utterance OR 70% of the time
        Suggest Gazing AWAY from user
FOR each RHEME
    If at end of utterance OR 73% of the time
        Suggest Gazing TOWARDS the user
```

Figure 6-5. Algorithm for controlling conversational gaze

Intonation Generator

The intonation generator implements three different strategies for controlling a Text-To-Speech (TTS) engine. The first strategy assigns accents and boundary tones based on a theme-rheme analysis, as described by (Prevost & Steedman, 1994) and shown in Figure 6-6. The intonation notation follows Pierrehumbert, in which L+H* and H* are high pitch accents, LH% and LL% are rising and low boundaries, respectively (Pierrehumbert & Hirschberg, 1990).

```

Within THEME:
    Suggest L+H* accent for NEW objects
    Suggest LH% boundary tone at end of THEME
Within RHEME:
    Suggest H* accent on NEW objects
    Suggest LL% boundary tone at end of RHEME

```

Figure 6-6. Algorithm for accent and boundary tone generation

The second intonation strategy suggests H* accents for all CONTRAST objects identified by the Tagger, following (Prevost & Steedman, 1994). The final intonation strategy simply suggests TTS pauses at CLAUSE boundaries.

Posture Shift Generator

Following a separate study on posture shifts in face-to-face conversation (Cassell, Nakano et al., 2001), I implemented a BEAT posture shift generator. In this study, we found that speakers tend to execute posture shifts—gross movements in limbs or stance—an order of magnitude more frequently at topic shift boundaries than at other times. Table 6-2 shows the primary results from this study. Posture shift frequencies were computed in two ways: per unit time (shifts/second) and per event (shifts/interval), with the latter computed only for transient events that lasted a brief time interval.

		Monologues		Dialogues	
		Shifts/second	Shifts/interval	Shifts/second	Shifts/interval
At	discourse	0.34	0.84	0.33	0.53
segment	boundaries				
Within	discourse	0.04		0.05	
segments					

Table 6-2. Posture Shift Frequencies Relative to Discourse Segments

The following BEAT rule was distilled from the frequency data for monologues (assuming a four-second duration per clause, the within discourse segment probability of a posture shift is 4 seconds x 0.04 shifts/second = 0.16):

```

FOR each CLAUSE
  If a TopicShift is indicated
    Then PostureShift 84% of the time
    Else PostureShift 16% of the time

```

Figure 6-7. Algorithm for Posture Shift Generation

The posture shift generator maintains posture state information for the character being animated, and outputs animation commands that move the character from its current posture into another one based on the animations available. Topic shifts are currently tagged in the Language Tagging module based on the occurrence of discourse markers in the input utterance (Schiffirin, 1987).

6.3.5 Behavior Selection

The Behavior Selection module analyzes the tree that now contains many, potentially incompatible, gesture suggestions, and reduces these suggestions down to the set that will actually be used in the animation. The selection process utilizes an extensible set of filters which are applied to the tree in turn, each of which can delete behavior suggestions which do not meet its criteria. In general, filters can reflect the personalities, affective state and energy level of characters by regulating how much nonverbal behavior they exhibit. Currently, two filter strategies are implemented: conflict resolution and priority threshold.

Conflict Resolution Filter

The conflict resolution filter detects all nonverbal behavior suggestion conflicts (those which physically cannot co-occur) and resolves the conflicts by deleting the suggestions with lower priorities. Conflicts are detected by determining, for each animation degree-of-freedom (DOF), the suggestions which co-occur and require that DOF, even if specified at different levels of the XML tree. For each pair of such conflicting suggestions (in decreasing order of priority) the one with lower priority is deleted unless the two can be co-articulated (e.g., a beat gesture on top of an iconic gesture). However, even in the case of co-articulation, two behaviors are not permitted to start using the same DOF at the same point in time. The types of nonverbal behaviors, their required DOFs, and co-articulation relationships are expressed in an XML file referenced by the filter. The filter operates as follows. For each DOF, the behaviors that use that DOF are considered in order of decreasing priority. For each behavior, a check is made to see if any other behavior that uses the DOF conflicts with it (overlaps in word indices when co-articulation is not allowed, or starts on the same word index when co-articulation is allowed). If a conflict exists, the lower priority behavior is removed from the tree. This operation is $O(N^2)$, where N is the maximum number of behaviors that use any given DOF (less than 10 for typical sentences).

Priority Threshold Filter

The priority threshold filter simply removes all behavior suggestions whose priority falls below a user-specified threshold.

6.3.6 Behavior Scheduling and Animation

The last module in the XML pipeline converts its input tree into a set of instructions that can be executed by an animation system, or edited by an animator prior to rendering. In general, there are two ways to achieve synchronization between a character animation subsystem and a subsystem for producing the character's speech (either through a TTS engine or from recorded audio samples). The first is to obtain estimates of word and phoneme timings and construct an animation schedule prior to execution (see Figure 6-8). The second approach is to assume the availability of real-time events from a TTS engine--generated while the TTS is actually

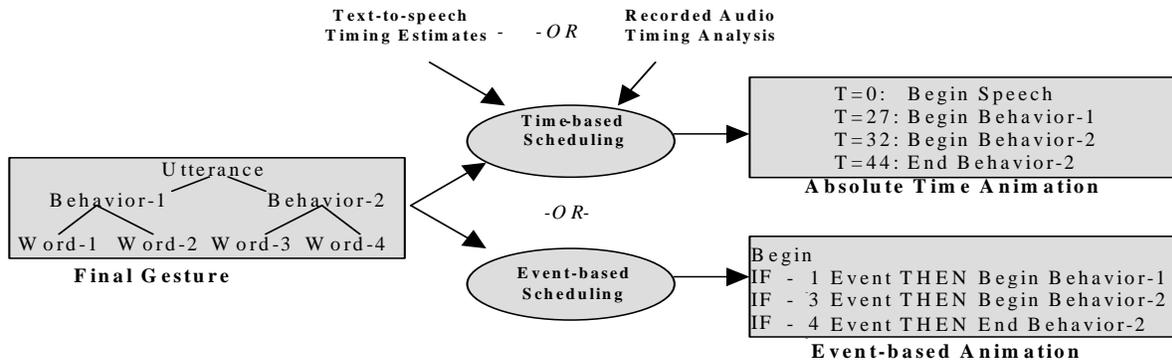


Figure 6-8. Scheduling Process

producing audio--and compile a set of event-triggered rules to govern the generation of the nonverbal behavior. The first approach must be used for recorded-audio-based animation or TTS engines such as Festival (Taylor, Black, & Caley, 1998), while the second must be used with TTS engines such as Microsoft's Whistler (Huang et al., 1996). The current toolkit is capable of producing both kinds of animation schedules, but the discussion here will focus on absolute-time-based scheduling with a TTS engine such as Festival.

The first step in time-based scheduling is to extract only the text and intonation commands from the XML tree, translate these into a format for the TTS engine, and issue a request for word and phoneme timings. In our implementation, the TTS runs as a separate process. Thus part of the scheduling can continue while these timings are being computed.

The next step in the scheduling process is to extract all of the (non-intonation) nonverbal behavior suggestions from the tree, translate them into an intermediate form of animation command, and order them by word index into a linear animation proto-schedule.

Once the word and phoneme timings become available, the proto-schedule can be instantiated by mapping the word indices into execution times (relative to the start of the schedule). The schedule can then also be augmented with facial animation commands to lip-sync the phonemes returned from the TTS engine. Figure 6-9. shows a fragment of an animation schedule at this stage of compilation.

```
<VISEME time=0.0 spec="A">
<GAZE word=1 time=0.0 spec=AWAY_FROM_HEARER>
<VISEME time=0.24 spec="E">
<VISEME time=0.314 spec="A">
<VISEME time=0.364 spec="TH">
<VISEME time=0.453 spec="E">
<GAZE word=3 time=0.517 spec=TOWARDS_HEARER>
<R_GESTURE_START word=3 time=0.517 spec=BEAT>
<EYEBROWS_START word=3 time=0.517>
```

Figure 6-9. Example Abstract Animation Schedule Fragment

The final stage of scheduling involves compiling the abstract animation schedule into a set of legal commands for whichever animation subsystem is being used. This final compilation step has also been modularized in the toolkit. In addition to simply translating commands it must concern itself with issues such as enabling, initializing and disabling different animation subsystem features, gesture approach, duration and relax times (the abstract schedule specifies only the peak time at start of phrase and the end of phrase relax time), and any time offsets between the speech production and animation subsystems.

6.3.7 Extensibility

BEAT has been designed to fit into existing animation systems, or to exist as a layer between lower-level expressive features of motion and higher-level specification of personality or emotion. To date, I have developed Translator modules for BEAT that generate “dope sheets” for professional animators to use to guide hand-drawn animations, for a humanoid animation system called Pantomime (developed at the MIT Media Lab), for the “light-weight” ECA described in chapter 8, and for the automatic production of Gantt charts such as the one in Figure 6-10. Other developers have implemented a Translator for Alias/Wavefront’s Maya animation system.

BEAT has also been designed to be extensible in several significant ways. First, new entries can easily be made in the knowledge base to add new hand gestures to correspond to domain object features and actions. Second, the range of nonverbal behaviors, and the strategies for generating them, can easily be modified by defining new behavior suggestion generators. Behavior suggestion filters can also be tailored to the behavior of a particular character in a particular situation, or to a particular animator’s style. Animation module compilers can be swapped in for different target animation subsystems. Finally, entire modules can be easily re-implemented (for example, as new techniques for text analysis become available) simply by adhering to the XML interfaces. Any kind of flexibility to the system derives from the ability to override the output from any of the modules simply by including appropriate tags in the original text input. For example, an animator could force a character to raise its eyebrows on a particular word simply by including the relevant EYEBROWS tag wrapped around the word in question. This tag will be passed through the Tagger, Generation and Selection modules and compiled into the appropriate animation commands by the Scheduler.



Figure 6-10. Gantt Chart

6.4 Extensions to BEAT for Relational Agents

Several extensions were made to BEAT to support the second-generation relational agent presented in chapter 8.

6.4.1 Attitudinal Functions

As discussed above, one of the most consistent findings in the area of interpersonal attitude is that immediacy behaviors—close conversational distance, direct body and facial orientation, forward lean, increased and direct gaze, smiling, pleasant facial expressions and facial animation in general, nodding, frequent gesturing and postural openness—demonstrate warmth and liking for one’s interlocutor and engagement in the conversation. BEAT was extended so that a subset of these cues (based on the nonverbal channels available) would be generated based on whether the ECA’s attitude towards the listener was relatively neutral or relatively warm.

Since BEAT is designed to over-generate, and produce nonverbal behaviors at every point in an utterance that is sanctioned by theory, attitudes are effected primarily by reducing the number of suggested nonverbal behaviors, as appropriate. For example, in a warm stance (high immediacy), fewer gaze away suggestions are generated, resulting in increased gaze at the interlocutor, whereas, in the neutral stance (low immediacy), fewer facial animation (eyebrow raises and headnods) and hand gesture suggestions are generated.

Cues that are encoded through relative frequency of behavior—increased gaze, facial animation, nodding and gesturing—are implemented by means of a StanceManager module which tracks the relational stance for the current utterance being processed, and is consulted by the relevant behavior generators at the time they consider suggesting a new behavior. Centralizing this function in a new module was important for coordination—since attitude (and emotion in general) affects all behaviors systemically. Modifications to baseline BEAT behavior were made at the generation stage rather than the filtering stage, since at least some of the behaviors of interest (e.g., eyebrow raises) are generated in pairs and it makes no sense to filter out a gaze away suggestion without also filtering out its accompanying gaze towards suggestion.

Relational stance affects not only whether certain nonverbal behaviors occur (i.e. their frequency), but the manner in which they occur. To handle this, the Translator module consults the StanceManager at animation compilation time to get a list of modifications that should be applied to the animation to encode manner (the “adverbs” of behavior). Currently, only proximity cues are implemented in this way, by simply mapping the current relational stance to a baseline proximity (camera shot) for the ECA, however, in general these modifications should be applied across the board to all aspects of nonverbal behavior and intonation (ultimately using some kind of animation blending, as in (Rose, Bodenheimer, & Cohen, 1998)).

The StanceManager tracks stance by an attribute in the root-level UTTERANCE tag that simply specifies what the relational stance is for the given utterance. For example:

```
<UTTERANCE STANCE="WARM">Hi there.</UTTERANCE>
```

The generators for gaze, gesture, headnods, and eyebrow movement consult the StanceManager at the time they are about to suggest their respective behaviors, and the

StanceManager tells them whether they can proceed with generation or not. Currently, lacking empirical data, the StanceManager reduces the frequency of filtered behaviors by 50%.

6.4.2 Framing Functions

As mentioned above, people clearly act differently when they are gossiping than when they are conducting a job interview, not only in the content of their speech but in their entire manner, with many of these “contextualization cues” encoded in intonation, facial expression and other nonverbal and paraverbal behavior. Since relational agents must be able to interact in not only task-oriented frames but a myriad of social frames, they must be able to clearly demarcate which frame of interaction they are in, in order to be natural and effective.

Contextualization cues (the nonverbal behaviors which mark frame shifts) are currently implemented in the StanceManager. Frames are marked in the input text using XML tags, such as the following:

```
<UTTERANCE><EMPATHY/>Sorry to hear that you're stressed  
out.</EMPATHY></UTTERANCE>
```

During compilation, the Translation module keeps track of the current frame and when it detects a change in frame it consults the StanceManager for the animation instructions which encode the requisite contextualization cues.

Currently there are four frames implemented—TASK, SOCIAL, EMPATHY, and ENCOURAGE—to support the relational agent described in chapter 8. These combine with the relational stance to yield a final set of modifications to behavior generation and animation modulation, as shown in Table 6-3. Figure 6-11 shows several examples of the effects of stance and frame on proximity and facial expression. For example, in the high immediacy, ENCOURAGE frame condition (lower left cell of Table 6-3) the agent is displayed in a medium shot (half way between a wide, full body shot and a close up shot), has a smiling facial expression, and does 50% fewer gaze aways than the default BEAT behavior (thereby spending more time looking at the user).

The relational stance and frame tags must be manually added to BEAT input strings for the above mechanisms to work. Section 8.6 discusses how these tags are computed for the health advisor agent.

6.5 Summary

One of the main lessons learned from the study described in chapter 5 is that, if a relational agent is to interact with users in simulated face-to-face conversation, then it is crucial that its nonverbal behaviors be appropriate to the current frame of interaction. This section has presented a software architecture that can be used to generate natural conversational nonverbal behaviors for a relational agent, to not only convey information and regulate the flow of conversation, but to convey attitude and conversational framing as well. This framework is used in a second-generation relational agent described in chapter 7 and evaluated in chapter 8.

Frame	Relational Stance	
	High Immediacy (Warm)	Low Immediacy (Neutral)
TASK	Proximity=0.2 Neutral facial expression Less frequent gaze aways	Proximity=0.0 Neutral facial expression Less frequent gestures Less frequent headnods Less frequent brow flashes
SOCIAL	Proximity=0.2 Smiling facial expression Less frequent gaze aways	Proximity=0.0 Smiling facial expression Less frequent gestures Less frequent headnods Less frequent brow flashes
EMPATHY	Proximity=1.0 Concerned facial expression Slower speech rate Less frequent gaze aways	Proximity=0.5 Concerned facial expression Slower speech rate Less frequent gestures Less frequent headnods Less frequent brow flashes
ENCOURAGE	Proximity=0.5 Smiling facial expression Less frequent gaze aways	Proximity=0.1 Smiling facial expression Less frequent gestures Less frequent headnods Less frequent brow flashes

Table 6-3. Effects of Stance and Frame on Nonverbal Behavior.
Frequencies are relative to baseline BEAT behavior. Proximity of 0.0 is a full body shot (most distant); 1.0 is a close up shot on the face.

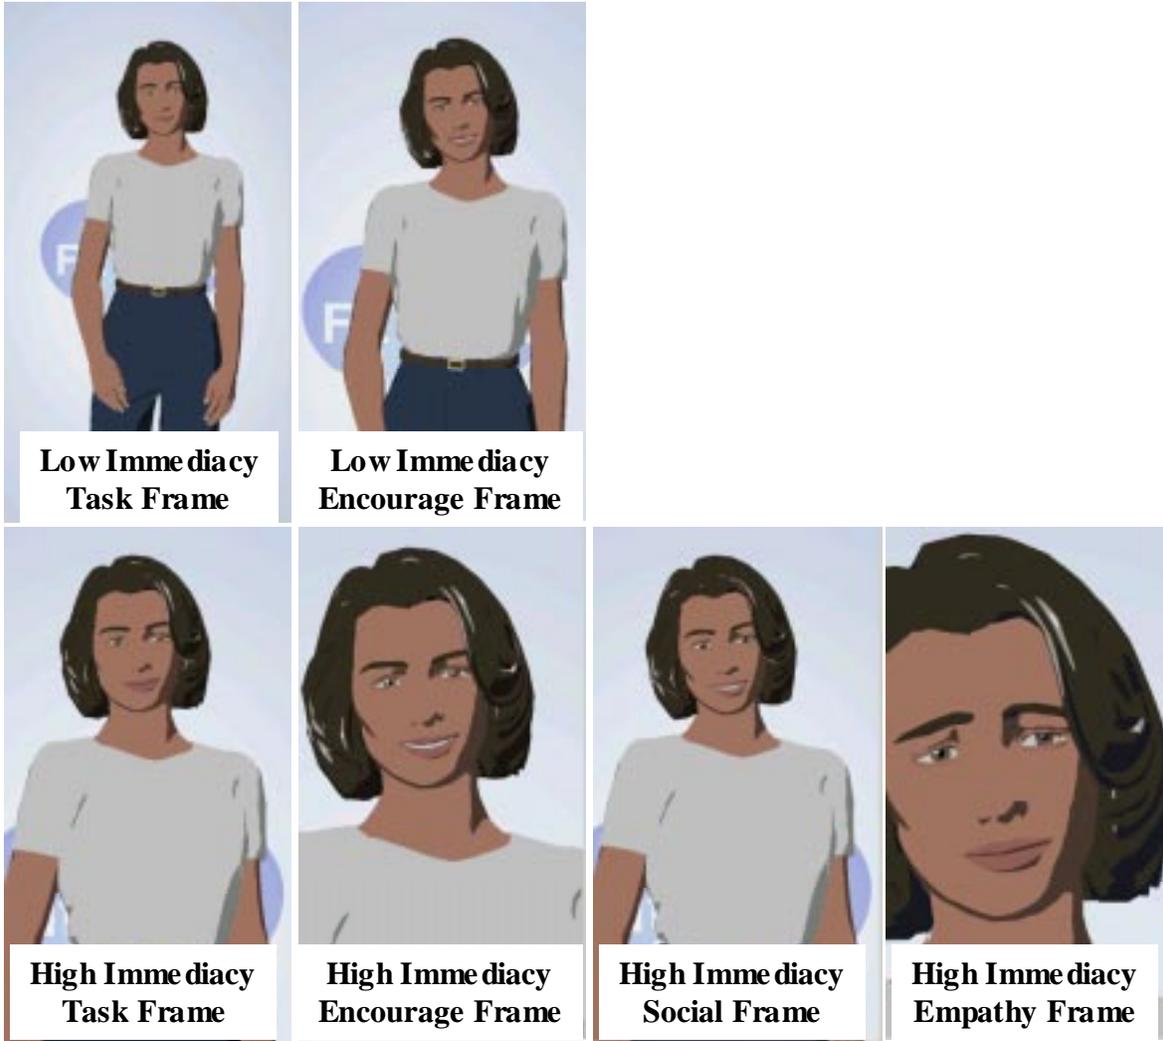


Figure 6-11. Example Effects of Stance and Frame on Proximity and Facial Expression

Chapter 7

LONG-TERM RELATIONSHIPS AND HELPING

The helping professions represent a potentially fruitful application area for relational agents who form long-term relationships with their users. There is a substantial empirical literature documenting a strong correlation between the quality of professional-client relationship and outcomes in psychotherapy, counseling and medical consultation. This chapter explores one such helping domain in detail—health behavior change for exercise adoption—and presents the results of two studies of professional exercise trainer-client interactions that are used to inform the design of a relational exercise advisor presented in the following chapter.

7.1 Introduction

As discussed in Chapter 1, there is a strong correlation between the quality of professional-client relationships and outcomes across a wide range of helping professions, but particularly in psychotherapy. The dimension of the therapist-patient relationship that is credited with the significant influence on outcome--the *working alliance*--is based on the trust and belief that the therapist and patient have in each other as team-members in achieving a desired outcome, and has been hypothesized to be the single common factor underlying the therapeutic benefit of therapies ranging from behavioral and cognitive therapies to psychodynamic therapy (Gelso & Hayes, 1998). As Gelso & Hayes, in their book on “The Psychotherapy Relationship”, put it:

The most fundamental component of the therapy relationship, we suggest, is the working alliance that develops between client and therapist. It is hard to imagine therapy being successful in the absence of a sound working alliance; indeed, empirical research supports the importance of the alliance to treatment outcome. Strong alliances are associated with more positive outcomes. ... The working alliance may be defined as “*the alignment or joining of the reasonable self or ego of the client and therapist’s analyzing or ‘therapizing’ self or ego for the purpose of the work.*” (Gelso & Hayes, 1998), pg. 9.

Attempts have been made to develop conversational, psychotherapeutic computer agents for significant disorders such as depression, but these have not received wide support because they lack the significant amount of natural language understanding and commonsense reasoning required for such extended therapy¹. However, behavior change is a sub-discipline of psychotherapy that has many well-understood, brief duration techniques, many of which have already been computerized in health behavior change applications such as smoking cessation and dieting (Celio, Winzelberg, Dev, & Taylor, 2002; Riva, Smigelski, & Friedman, 2000; Velicer & Prochaska, 1999). Empirical investigations have demonstrated that the

¹ See, for example, Sherry Turkle’s discussion of computer psychotherapy in (Turkle, 1995).

positive impact of a strong working alliance also holds for some areas of health behavior change, including treatment for drug and alcohol abuse.

The working alliance and its impact on health behavior change outcomes is thus an ideal application area for relational agents, because: 1) the working alliance is a well-understood and measurable type of relationship (e.g., (Horvath & Greenberg, 1989)); 2) health behavior change involves interventions over extended periods of time, and this provides a framework within which relational agents can establish persistent, long-term relationships; and 3) effective health behavior change is of direct benefit to subjects.

In order to evaluate the effectiveness of relational agents in health behavior change, I developed an exercise advisor agent that promotes exercise among MIT students. According to (Pinto, Cherico, Szymanski, & Marcus, 1998): 1) college students' participation in regular exercise offers both physiological and psychological benefits, having been shown to decrease anxiety and depression, reduce test anxiety, and improve self-esteem; 2) *Healthy People 2000* specifically identified post-secondary educational institutions as settings where young adults should be targeted for exercise promotion; and 3) surveys of college students' health habits indicate that only about 35% to 37% report having a regular schedule of physical activity.

The target behavior the exercise advisor agent helps users work towards is 30 minutes or more of moderate or vigorous activity on most days of the week. This is the latest recommendation from the Centers for Disease Control and Prevention and the American College of Sports Medicine (Pate et al., 1995) (the previous recommendation was for 20 minutes of vigorous activity three times per week). The daily recommendation represents an ideal framework within which to evaluate a relational agent, since it affords the opportunity for daily check-ins and chats with a relational health advisor agent over the duration of the intervention.

The exercise advisor agent will recommend walking as the form of exercise that users should adopt, even if they are already engaged in an exercise or sport. Walking is a low-risk activity that requires minimal equipment and can be performed almost anytime or anywhere (Jonas, 1995), plus it is a good habit for college students to get into since it will last them the rest of their lives, something that cannot be said of many intramural sports such as water polo or roller hockey.

Developing an autonomous exercise advisor agent that is capable of building relationships with people *and* influencing their exercise behavior requires knowledge spanning a number of disciplines. In the following sections I briefly outline prior work in the working alliance and exercise behavior change, as well as two pilot ethnomethodological studies of interactions between fitness trainers and their clients, all of which were drawn upon to inform the design of the exercise advisor agent (described in the following chapter).

7.2 Building the Working Alliance

The Working Alliance construct has three sub-components: a goal component, reflecting the degree to which the therapist and client agree on the goals of the therapy; a task component, reflecting the degree to which the therapist and client agree on the therapeutic tasks to be performed; and a bond component, reflecting the trusting, empathetic relationship between the client and therapist (Gelso & Hayes, 1998; Horvath & Greenberg, 1989).

Trust, rapport and empathy (reflected in the bond dimension of the alliance) is frequently mentioned in the literature on helping to be pre-requisites for successful therapy of any kind.

According to Okun, there are two stages to the helping process, with the first stage focused on building rapport and trust between the helper and helpee, in which the helper supports self-disclosure to uncover and explore as much information and as many feelings as possible, before moving on to the second stage of strategy planning, implementation and evaluation (Okun, 1997). In the literature on the working alliance, the relational factor that is most often mentioned as crucial in forming and maintaining the alliance is the patient's perception of the therapist's empathy for them (Gelso & Hayes, 1998).

Even in physician-patient interactions, physician empathy for a patient plays a significant role in prescription compliance, and a physician's *lack* of empathy for a patient is the single most frequent source of complaints (Frankel, 1995). Based on a meta-analysis of several studies, Buller and Street recommend establishing a "positive relationship" with the client before exercising authority (e.g., giving prescription) (Buller & Street, 1992). They also found that physicians who are more expressive, particularly of positive emotion, are more "satisfying" to their patients; that inconsistent or confused emotional expressions by the physician lead to more negative evaluations of them; and that expression of negative affect by physicians (anger & anxiety) were associated with higher compliance and better patient health (Buller & Street, 1992).

Empathetic language involves seeing the world through the client's eyes, with statements such as "It is awful." or "How awful it is." (Havens, 1986). Simple empathetic statements include: paraverbals (from backchannel feedback to crying out); adjectives of empathy (e.g., "awful", "wonderful"); accented adjectives ("How awful!"); and translations ("It is terrifying.", "Isn't it terrifying?"). Empathetic statements can also be more complex, such as "No wonder you were frightened!".

In addition to conveying an appropriate amount of empathy, trust and rapport can be developed through social dialogue (Bickmore & Cassell, 2001; Laver, 1975; Malinowski, 1923; Schneider, 1988), reciprocal deepening self-disclosure (though there are mixed opinions about how much self-disclosure a helper should do) (Altman & Taylor, 1973; Moon, 1998), emphasizing commonalities (Gill, Christensen, & Fincham, 1999), humor (Cole & Bradac, 1996; McGuire, 1994; Stafford & Canary, 1991), meta-relational communication (particularly emotional aspects) (Dainton & Stafford, 1993; Stafford & Canary, 1991), and expressing happiness to see the helpee (Okun, 1997).

Other relational strategies that could be used to effect increases in the bond dimension of the alliance include: talking about the past and future together (Planalp, 1993; Planalp & Benson, 1992), continuity behaviors (appropriate greetings and farewells and talk about the time spent apart) (Gilbertson, Dindia, & Allen, 1998), and reference to mutual knowledge. Specific language constructs that may also be effective include the use of inclusive pronouns (Havens, 1986), and politeness strategies (Brown & Levinson, 1987; Lim) and greeting and farewell rituals (Laver, 1981) indicative of a close relationship.

There is much less written on what a helper can do to increase the task and goal dimensions of the alliance. However, some of the strategies mentioned include: clarifying mutual goals as much as possible initially, and frequently reminding the helpee of these goals; clearly motivating any tasks the helpee is asked to perform; and being clear about the roles and expectations that the helper and helpee will have (a form of meta-relational conversation) (Okun, 1997).

7.2.1 Scheduling of Techniques

In psychotherapy, the working alliance seems to quickly reach a peak (after 3-7 sessions) then level off (Horvath & Symonds, 1991). This indicates that most relational techniques should be in use by the exercise advisor by the seventh session. In general, the strategies that should be varied over time include:

- Increasing breadth of topics (mostly during social dialogue).
- Increasing depth of topics (mostly during social dialogue).
- Increasing amount of small talk (and thus overall length of interaction).
- Increasing amount of information the agent knows about the user (enables the above 3, plus telling the agent a lot information is a type of "investment" made in the relationship).
- Increasing use of empathy, agreement with the user, and humor (demonstrates increased knowledge of subject, and that agent and subject have a "shared meaning system" (Duck, 1994)).
- Initial talk about the relationship to ensure there are no initial problems, then occasional check-in or as needed.
- References to past interactions and mutual knowledge should increase over time as the agent learns more about the subject.
- References to future interactions, inclusive pronouns, expressing happiness to see the user, nonverbal immediacy behaviors, and the use of greeting routines, forms of address and politeness strategies indicative of a close helping relationship can be used continuously from the first interaction.

7.3 Exercise Behavior Change

One of the most widely used approaches to health behavior change is the Transtheoretical Model (Prochaska & Marcus, 1994), which defines a set of five stages that people go through when effecting change, along with a set of measures and therapeutic change processes effective within each stage. The stages are: precontemplation, when people have no intention of changing their behavior; contemplation, when people intend the change within the next six months; preparation, when individuals intend to take action within the next month; action, when they have done something to change their behavior within the past six months; and maintenance, when they have maintained the desired change for at least six months. In addition to assessing stage of change directly, two key constructs that are typically assessed in studies on this model are decisional balance and self-efficacy. Decisional balance reflects the pros vs. cons of changing behavior (based on theories of human decision-making), and is typically seen to change from more cons than pros in precontemplation and contemplation to the opposite situation in action and maintenance. Self-efficacy, in this case people's degree of confidence that they can exercise, is widely cited as the single most significant predictor of exercise behavior change (Dishman, 1994), and is also a measure of progress, especially in action and maintenance stages. Because individuals in precontemplation and maintenance typically require very different types of interventions, I decided to focus on behavior change strategies for people who were in the contemplation, preparation or action stages of change only.

Of the many exercise behavior change strategies described in the literature (e.g., (Jonas, 1995; King, 2001; Knapp, 1988)), the following are among the most frequently mentioned:

- **Positive Reinforcement** - Positive reinforcement is any stimulus or event that is presented during or following the behavior that increases the future rate of the behavior upon which it is made contingent. To be maximally effective, reinforcement should be delivered during or shortly after the behavior. For most beginners, social support and attention (e.g., coach praise) during and after exercise reinforces the exercise habit (the exceptions being those individuals who do not like to exercise with others). Reinforcement should be provided as frequently as possible in the early stages of exercise habit formation.
- **Stimulus Control** - This includes the use of cues and prompts that stimulate, or are reliably followed by exercise. Examples include: teaching people to lay out their exercise clothes the night before they exercise; spending time with others who exercise frequently; or teaching exercisers to encourage others to ask about their exercise. The removal of cues for competing behaviors is also important (e.g., exercising first thing in the morning when there are few competing demands on one's time).
- **Goal Setting** - The achievement of one's exercise goals has a significant impact on adherence to exercise. Near-term or daily goals should be set by individuals, because individuals who believe that their choices help determine their exercise prescription have better adherence. Time-based goals work better overall than distance-based goals.
- **Shaping** - This is the process by which behavior is broken down into a series of successive approximations that are gradually progressed to the desired behavior. The most common mistake made by many beginners is initiating exercise programs at too high an intensity, frequency, and duration. Thus, the first rule of shaping is to start with very simple, easily performed behaviors and work up. The primary goal during the first 8-12 weeks should be to establish the habit of exercise and not physical conditioning.
- **Self-Monitoring** - Self-monitoring involves individuals keeping records of their own exercise behavior over time, ideally in a visual format (e.g., a plot of minutes of exercise per day) posted in a highly visible location. This strategy is not only good for beginners, but for maintainers as well.
- **Dissociative Strategies** - New exercisers should be taught to not focus on their body when exercising (e.g., pay attention to the scenery when walking rather than on how much their feet hurt).
- **Coping Thoughts** - Subjects should be taught that when they become aware of negative thoughts or images, such as "I'm not doing as well as I should...I have a long way to go...This is boring", they are instructed to make positive self-statements such as "I'm doing well to exercise at all today...I'm nearly half-way finished...Let me notice what's going on around me."
- **Education** - Information about the benefits of exercise, as well as proper exercise techniques, can influence individuals' decisions to exercise.

- **Decisional Balance** - Simply going through the exercise of thinking about the pros and cons of a beneficial behavior such as exercise can influence individuals' decision about engaging in the behavior or not.
- **Social Support** - For those individuals who enjoy working out with someone, a "workout buddy" can be a great source of mutual social support, positive reinforcement (from both praise and the pleasure of spending time together), and an aide to cognitive dissociation (distraction from discomfort and boredom).
- **Relapse prevention training** - Of interest primarily to maintainers, this set of techniques includes: 1) teaching individuals that lapses are normal and likely; 2) cautioning against the "abstinence violation effect" (the belief that once any rule about the behavior is broken, then total relapse is inevitable), by teaching individuals that adherence is a continuous--not dichotomous--variable; 3) instructing individuals to avoid high-risk situations (e.g., starting to watch TV at exercise time); 4) guiding individuals through a "planned relapse" in which the educator/coach can provide support while they learn the techniques necessary for recovering from a lapse.

In addition to these, several additional strategies can be borrowed from cognitive-behavioral psychotherapy and applied to exercise behavior change:

- **Reasons / Affirmations** - Individuals should be asked why they want to exercise or get in shape, in as specific language as possible. These reasons can then be used to help motivate them when they are having problems (e.g., "Remember, you wanted to exercise so that you could feel better about your body.") and they can also be instructed to repeat these statements to themselves frequently ("I'm exercising so that I can feel better about my body.").
- **Reasons Not / Counter Arguments** - People come up with a wide range of excuses for not exercising, but these should be met with counter arguments that the coach can tell them or have them tell themselves (e.g., "I know I feel that I should study, but getting exercise will actually help me think more clearly and be more productive."). Obstacles to exercise can be categorized and dealt with accordingly. In one study of high school students, the most frequent perceived barriers to exercise included: lack of time, bad mood, lack of energy, lack of self-discipline, discomfort, cost, thinking that exercise is not fun, feeling too self-conscious, feeling too stressed, lack of support from family and friends, and illness and injury (Allison, Dwyer, & Makin, 1999).
- **Performatives** - Performatives are explicit statements the coach makes about their belief in the capabilities of their clients to achieve change (e.g., "I really believe you can do it.") (Havens, 1986).

7.4 Ethnomethodological Studies of Trainer-Client Interactions

In order to understand specific techniques and language used by professional exercise trainers, two ethnomethodological studies of interactions between personal exercise trainers and their clients were conducted to determine the range of conversational, relational and nonverbal strategies typically used in these interactions. The first study involved five face-to-face meetings between a personal trainer and college undergraduate clients, which were videotaped for subsequent analysis. The second study involved five students holding daily text chat meetings with a trainer over a period of two weeks. The second study was designed

to simulate the interaction protocol that would be used in the exercise advisor system, with subjects recording both daily pedometer readings and minutes of physical activity, and the trainer being able to access these records during the daily chat sessions.

7.4.1 Pilot 1: Half-hour face-to-face interactions

In this pilot study, three MIT students had 30-minute introductory meetings with a fitness trainer to discuss their fitness goals and needs. Two of the three students returned a week later for second 30-minute follow-up meetings. Summaries of the topics discussed by one subject on her first and second meeting are presented below in Tables 7-1 and 7-2, along with an analysis comparing surface features of these two interactions in Table 7-3.

Although the trainer spent over half of her time discussing anaerobic exercise and nutrition—topics that would not be covered by the exercise advisor agent—there were a number of important techniques learned from these pilots.

The initial session with all clients took an interview format, in which the trainer spent most of her time asking questions then making recommendations, with only a few instances in which she opened the floor to client questions. Follow-up sessions were much less structured in nature, with the trainer following up on the client’s performance and making many recommendations, but with much of the discussion driven by specific needs or questions presented by the client.

Initial sessions started off with the trainer asking clients what their long-term (six-month) and short-term (one-month) goals were for exercising, and what kinds of benefits they expected to get out of establishing an exercise routine. She would then find out if they had exercised or engaged in sports in the past, and if they were currently engaged in any sport activity, so that she could tailor her recommendations accordingly. She would also find out if they had any health problems that would prevent them from exercising or restrict the kind of activities they could engage in. Following this, she would make specific recommendations for an exercise routine that the client could get started on, plus provide a range of tips and instructions for performing specific exercises and staying motivated.

Turns	Segment / Topic (Duration - 40:27)
1-6	Greeting
7-13	Client's Top-level Goal: Why Exercise?
14-18	Client's Past History of Exercising
19-59	Medical Problems
60-136	Client's Goals
137-155	Place to Exercise
156-294	Overcoming Obstacles to Exercising
295-625	Strength Training
625-659	Scheduling of Exercise Routine
660-699	Tips
700-734	Open Floor to Client Questions
735-744	Closing/Farewell

Table 7-1. Summary of Exercise Advisor Pilot 1, Subject 1, First Meeting

Turns	Segment	(Duration – 32:09)
1-3	Greeting	
4-8	Smalltalk	
9-40	Check-in: How did everything go?	
41-118	Check-in: How did strength training go?	
119-352	Open Floor: Any questions about particular exercises?	
353-371	Gym Schedule	
372-430	Stretching	
431-467	Exercise Routine	
468-517	Daily Routine/Schedule	
518-544	Cardio	
545-556	Daily Routine/Schedule	
557-558	Open Floor - Any Questions?	
559-593	Review	
594-603	Diet/Nutrition	
604-624	Exercise Routine	
625-634	Open Floor - Any problems?	
635-647	Schedule / Daily Routine	
648-673	Coping with Holiday Meals	
674-702	Small talk	
703-708	Scheduling	
709-722	Self-Monitoring	
723-754	Self-Motivation/Reward	
755-771	Obstacles	
772-782	Closing / Farewell	

Table 7-2. Exercise Advisor Pilot 1, Subject 1, Follow Up Meeting

Follow up meetings began with the trainer finding out how the client had been doing during the previous week on their exercise program and determining if they had had any specific problems that needed to be addressed. Following this, the meetings were very unstructured with the trainer answering questions and making recommendations in response to client needs or questions.

A nonverbal behavior of particular interest was the range of facial expressions used by the trainer to mark affective state and for contextualization cues². Three expressions regularly stood out: a neutral face during the majority of the interactions, and in particular during information exchanges; a concerned face during empathy exchanges (e.g., following the client disclosing that she was allergic to seafood); and a happy/smiling face during greetings, farewells, social dialogue and humor.

² Gumperz defined “contextualization cue” to be any aspect of the surface form of utterances which can be shown to be functional in the signaling of interpretative frames (Gumperz, 1977)

	First Meet	Followup Meet
Duration	0:40:27	0:32:09
Total discourse segments / topics	12	24
Total turns	744	782
Trainer turns	382	391
	51%	50%
Client turns	362	391
	49%	50%
Trainer backchannel-only turns	69	75
Client backchannel-only turns	218	184
Total content turns	457	523
Trainer content turns	313	316
	68%	60%
Client content turns	144	207
	32%	40%
Trainer laughs	7	23
Client laughs	35	62
Small talk turns (completely unrelated to exercise)	2	35
Small talk turns (related but non-instrumental)	21	78
Total words	8294	7432
Trainer words	6262	5037
	76%	68%
Client words	2032	2395
	24%	32%

Table 7-3. Exercise Advisor Pilot 1, Subject 1, Comparative Analysis

Regarding relational dialogue, there was almost no small talk in the first meeting but 35-78 turns of it in the follow up meeting (depending on how these turns are classified). This was accompanied by significantly more laughter in the follow up meeting; 85 instances vs. 42 instances in the first meeting. The trainer also made very frequent use of empathetic feedback, especially when the client mentioned problems that she was having (e.g., client: “I mainly have problems with cold air and stuff.”; trainer: “Cold? Yea, I’m the same way. ... I have asthma too. ... So I feel the same way.”). Unfortunately, since each client only had two sessions with the trainer, there was not much opportunity for them to develop much of a relationship over time, and thus many of the relational strategies mentioned in section 7.2 were not observed to occur in the dialogues.

7.4.2 Pilot 2: Ten-minute text-chat interactions

A second pilot study was conducted in which MIT students interacted with a fitness trainer via a text chat system, intended to parallel the protocol that would be used by the automated exercise advisor. In order to provide an objective measure of physical activity, subjects wore pedometers and recorded their pedometer readings and estimates of time in different categories of physical activity (following (Sallis, 1997)) using web forms. The

subjects then participated in a 5-10 minute text chat session with a fitness trainer. During these sessions, the subjects and the trainer could view the subject's exercise data (see Figure 7-1). A total of five students participated in the study over a period of 3 weeks. Due to numerous technical and logistical problems only 13 usable chat logs were collected, averaging 50 turns each in length. Figure 7-2 shows an example of a typical chat session.

The chat logs were analyzed with the fitness trainer at the conclusion of the study in an attempt to understand her goals and strategies used during the interactions. Her general strategy was to first find out the current physical and mental state of the client (e.g., by asking "How are you?" "How is it going?"). She would next determine if the client had exercised since they last talked, in particular following up on any specific commitment to exercise made by the client in the prior chat. She would always give congratulations and praise if they met their commitment, and would be especially emphatic if the client had to overcome some obstacle in order to meet their commitment. She would then follow up on any problems the client was having (e.g., getting new exercise shoes or seeing a physician about an injury), and would be persistent in making sure these were resolved, so that they would not become obstacles to exercising. Finally, she would get the client to make a specific commitment to exercise before the next time they chatted. This often took several turns of badgering and

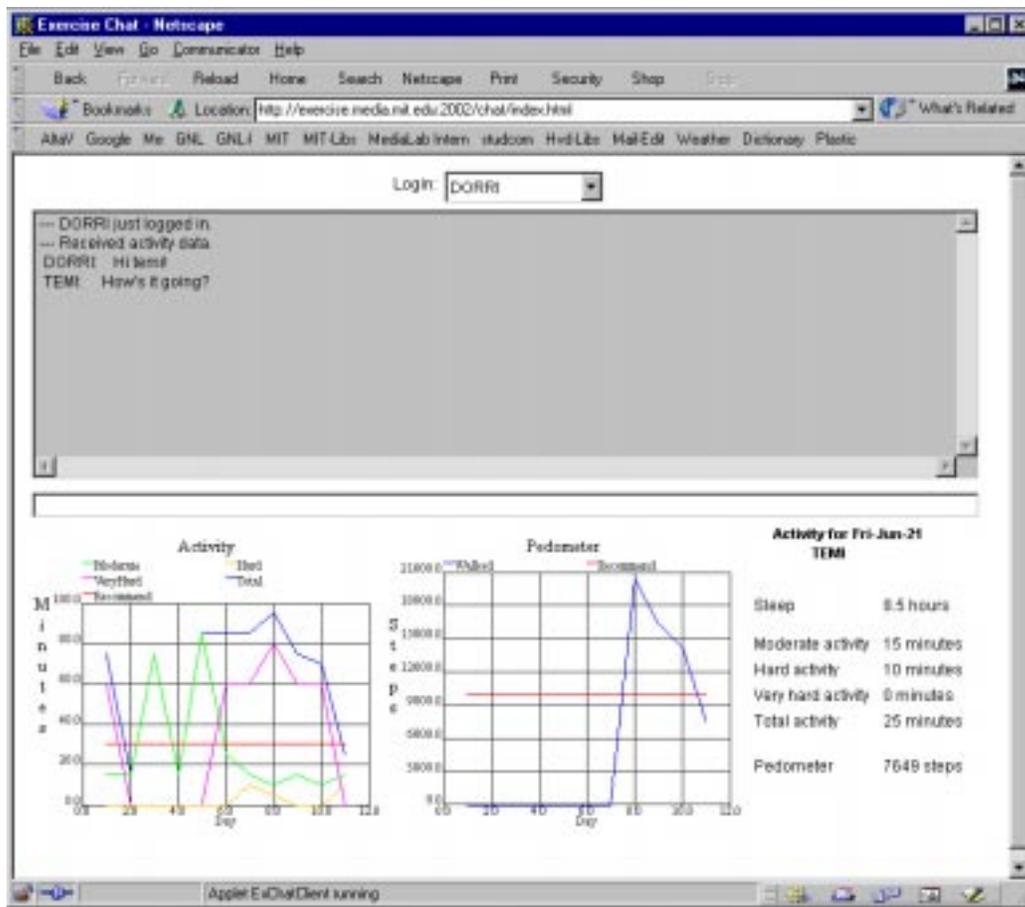


Figure 7-1. Chat System Interface Developed for Pilot 2

1. CLIENT: hi
2. TRAINER: Hi Mary How are you doing today?
3. CLIENT: good - how are you?
4. TRAINER: fine thanks. How was your workout yesterday?
5. CLIENT: it was alright, but I'm starting to get tired of it already
6. TRAINER: hmmm...what can we do?
7. CLIENT: there's a free ballroom dancing class tonight - so I'm going to try that
8. CLIENT: I think it should be more fun
9. TRAINER: good idea, brilliant actually
10. TRAINER: it's helpful to vary your routine
11. TRAINER: back to getting tired of your workout, where are you training?
12. CLIENT: weight room in my dorm
13. CLIENT: it's very small
14. TRAINER: many people use it?
15. CLIENT: about 3 people at a time so far
16. TRAINER: sometimes its helpful to have more people around , more energy happening the ballroom dancing should be fun
17. TRAINER: going with friends?
18. CLIENT: yup
19. TRAINER: that's nice
20. CLIENT: the weight room gets kind of boring i guess
21. TRAINER: have you thought about setting any fitness goals for yourself?
22. CLIENT: i think i want to walk around more - explore boston
23. TRAINER: that would get your walking in Boston is a good walking city
24. TRAINER: have you been here long?
25. CLIENT: just since fall
26. TRAINER: like it so far?
27. CLIENT: yup - actually, are there any places you would recommend walking to?
28. TRAINER: Newbury Street, the north end, Chinatown if you like Chinese food, Quincy Market for the tourist scene
29. CLIENT: what about parks (besides boston commons and the gardens next to it)?
30. TRAINER: you can walk along the waterfront in the north end. south boston ,which would be a trek, has a long beach that you can walk
31. CLIENT: ok thanks – i'll try those sometime this summer
32. TRAINER: your welcome it should be fun
33. TRAINER: are you going to workout today ?
34. CLIENT: ballroom dancing...
35. TRAINER: right, short memory sorry shall we touch base tomorrow?
36. CLIENT: same time?
37. TRAINER: yup
38. CLIENT: alright – talk to you then
39. TRAINER: have fun tonight Mary
40. CLIENT: thanks

Figure 7-2. Exercise Advisor Pilot 2 Sample Chat Session

occasionally the trainer would need to lower or change her expectations, by asking the client what they would like to do for exercise, and letting anything be better than nothing.

The trainer would also give tips, acquire information about the client's preferences (e.g., whether they like to exercise alone or with friends), and occasionally joke as the opportunity arose.

Many of the relational and behavior change strategies discussed in sections 8.2 and 8.3 were used by the trainer in these sessions, including:

- Forms of address in greeting and parting. "Hi Mary!", "ok Tom nice chatting with you goodbye!"
- Use of inclusive pronouns: "hmmm...what can we do?"
- Talking about the past and future: "...shall we touch base tomorrow?"
- Continuity behaviors (asking about what client did since they last talked): "How was your workout yesterday?"
- Social dialogue: client, "i think i want to walk around more - explore boston" ... trainer, "have you been here long?" client, "just since fall." trainer, "like it so far?" etc.
- Reference to mutual knowledge: "hi mary how was the ballroom dance class?"
- Performatives: "well i'm rooting for you."
- Role clarification: "my role is to encourage you to exercise and to answer any questions you might have."
- Empathy: client, "my toe hurts a lot" ... trainer, "ouch it definitely would be good to get that looked at".

7.5 Summary

This chapter presented a range of therapeutic and relational techniques that should be incorporated into a relational exercise advisor agent that is to be effective at changing clients' exercise behavior.

Relative to the theory presented in chapter 3, the working alliance can be seen as the client's trust in the benevolence and competence of the helper in working with them to achieve change. Establishing and maintaining these expectations is relational work, and relies on all of the techniques described above, although empathy on the part of the helper seems to play a particularly important role.

The following chapters present the design and evaluation of a relational agent that plays the role of an exercise advisor who attempts to help users increase their level of physical activity over an extended period of time. A great deal of information from the two pilot studies was incorporated into the design of this agent, from the overall structure of the dialogues used to the actual verbatim use of utterances made by the trainers:

- The structure observed in the first interactions of Pilot 1 was used as a template for the first two dialogues for the exercise advisor agent, in which the agent spends most of the time discussing clients' exercise goals and telling them what they can expect in future interactions. Subsequent agent dialogues take the form of typical interactions in Pilot 2, following the general template below (with significant daily variations):
 1. Greeting
 2. Determine the physical and emotional state of the user.

3. Social dialogue.
 4. Follow up on the user's commitment to exercise made in the previous interaction.
 5. Offer exercise tips.
 6. Get the user to commit to a specific amount of exercise before the next interaction.
 7. Farewell
- The nonverbal behavior of the trainer observed in Pilot 1 was used to guide the design of the conversational frames and facial expressions used for the exercise advisor agent (described in chapter 6).
 - Finally, all of the relational and behavior change strategies discussed above (section 7.4) were implemented in the exercise advisor dialogues.

Chapter 8

DESIGN OF A RELATIONAL AGENT FOR LONG-TERM HELPING

The development of relational agents for long-term use presents many significant design and engineering challenges that have yet to be addressed by systems that are intended for single session interactions. Some of these challenges include: persistence (maintaining state across sessions); continuity (taking the context of past interactions into account in the current interaction); trajectories of change (how the system should change over time to meet the user's expectations of change for the social role the agent is playing); and variability (how the system should change its behavior in response to the current interactional context). This chapter presents the design of a relational agent that addresses many of these challenges.

Although deep models of relational behavior are required for maximum generality and adaptability of relational agents, the agent presented in this chapter does not rely on an explicit model of its relationship with the user, aside from simply keeping track of the number of interactions they have had together and an accumulation of propositional information about the user gleaned from prior conversations. The interactions are largely scripted, but take into account expected changes in relational behavior over time and rely on accumulation of shared knowledge (common ground).

Two versions of this agent were created in parallel in order to support the experiment described in the following chapter: a “non-relational” version and a “relational” version. The non-relational version carried out all of the task dialogue of the relational version, but with all of the purely relational behaviors (verbal and nonverbal) ablated.

8.1 Requirements and Specifications

The exercise advisor system—known as “FitTrack”—was designed to increase physical activity in free-living individuals (primarily, but not exclusively MIT students), moving them from below 30 minutes a day of moderate-or-better activity up to 30 minutes a day or more of moderate-or-better physical activity over a one-month period of time. In order to support free-living subjects, the system had to be designed to run on typical home computers so that participants could interact with the system on a daily basis to input their day's activity and have a conversation with the animated exercise advisor. This requirement alone presented significant challenges, since most embodied conversational agents are designed to run in a laboratory setting on high-end graphics workstations and are not built to run reliably without the continuous support and intervention of the developers. One of the major implications of this requirement is that a new, light-weight animation system needed to be developed that could run in real-time, synchronous with a text-to-speech engine.

Text-to-speech was selected over recorded speech, due to the very large number of utterances that needed to be produced over the life span of the interaction, and due to the variability in utterances required to support many relational strategies (e.g., referring back to prior conversations or using the user's name in greetings). Text-to-speech was selected over “text balloon” display, since there is no way to effectively convey synchronization between

ECA nonverbal behaviors and text, and the timing of nonverbal behavior itself carries important meaning.

Dynamic menus were used as the primary means of input (as in (Rich, Sidner, & Lesh, 2001)) rather than speech recognition or free text natural language input, since both of these modalities lead to failed expectations of unconstrained natural language understanding, which, in the case of conversational speech is compounded by high recognition error rates.

Another major design decision was that the system would have a client-server architecture, with the client (including the ECA) designed to be as simple, reliable and light weight as possible (with respect to user system resources), with all logic and dialogue handling performed on the server, and all data (both exercise information and dialogue state) stored on the server. This enabled the inevitable changes and bug fixes to be performed on the server without distributing new versions of the software to users, allowed user data to be kept in a secure location, and allowed the system to be used on user computers with a wide range of processor speeds and without graphics acceleration.

The exercise advisor system was also designed to support an evaluation experiment, described in full in chapter 9. Briefly, this experiment involves subjects interacting with the system daily for a one-month period of time. Each session is designed to last 10-15 minutes, during which subjects enter their physical activity (time estimates and pedometer readings) into a web form, occasionally fill out additional self-report questionnaires, view their progress (self-monitoring), have a brief conversation with the exercise advisor agent, and are finally presented with a daily page of educational content about exercise. There are three treatments

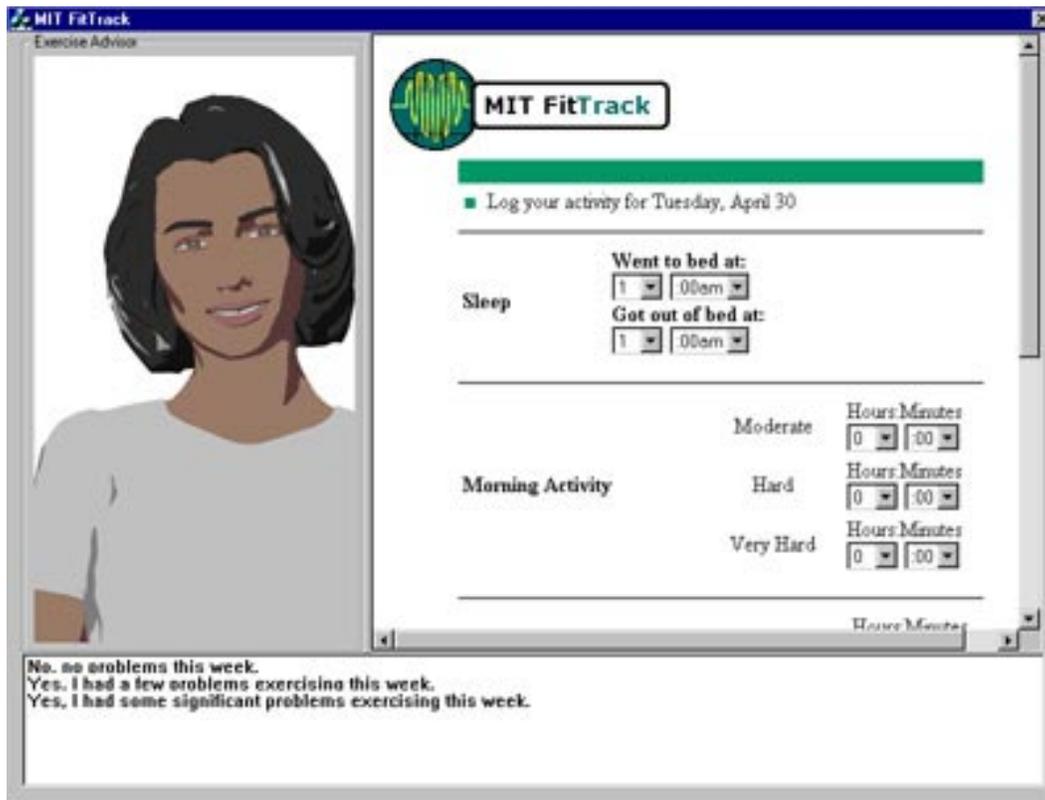


Figure 8-1. FitTrack Client Software Interface

in the study: the RELATIONAL group has access to all system functionality, including interactions with an agent that uses relational strategies to attempt to build a working alliance with them; the NON-RELATIONAL group has access to all system functionality, except that all purely relational strategies (verbal and nonverbal) have been ablated from the agent interactions; and the CONTROL group has access to all system functionality except that they do not get to interact with the advisor agent (nor have any knowledge of its existence). These requirements led to a system design in which the client application consisted of a web browser--used to deliver forms and educational content--coupled with the exercise advisor agent, and in which the agent could be hidden by a second browser window for the CONTROL group. Figure 8-1 shows a screen shot of the final client application with the agent displayed. The largest window pane in the interface is the “content” pane (used to display forms, educational content and self-monitoring charts), the left-hand pane is the “help” pane (alternates between a second browser window for help pages and the ECA animation), and the bottom pane is the “input” pane (used for user input during ECA dialogues; blank at all other times).

8.2 System Architecture

Figure 8-2 shows the overall architecture of the exercise advisor system. When started, each client connects to the dialogue server via a persistent TCP/IP connection, and connects intermittently to the web server using standard world-wide-web protocols (HTTP transactions over TCP/IP).

The relational database on the server is used to store all information about users and their sessions with the system, including all messages passed between the server and clients, in order to support debugging and usage analyses unanticipated at the time the system was developed. An overview of the database schema is given in Table 8-1.

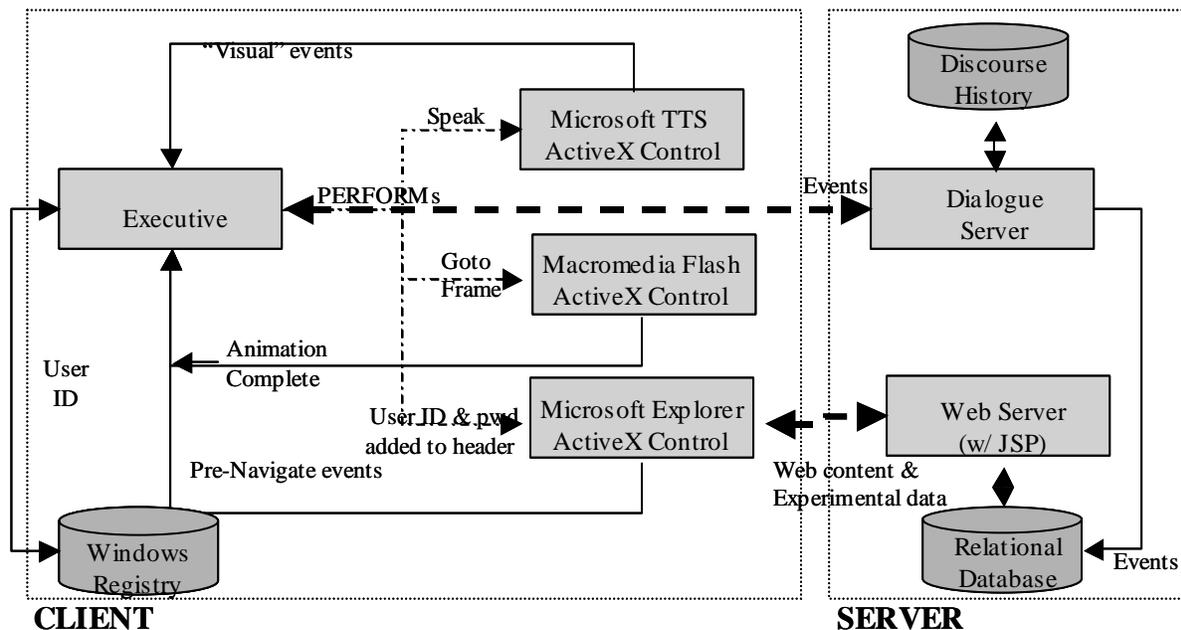


Figure 8-2. FitTrack System Architecture

Table	Description
Users	Information about individual subjects in the experiment, including USER ID, password, email address, day of the study, given name and experiment state information.
ActivityLogs	Subjects' estimates of time in different intensity categories of physical activity as well as their pedometer readings for each day.
Sessions	Records of each individual client session from login through termination.
Events	All events that occurred for a given session with a client, including client server messages, dialogue state changes, and internal server events.
Surveys	Records describing each questionnaire filled out by a subject (survey type, survey instance, user and session).
Responses	Individual responses to survey questions.
Properties	Information about the user that is kept between sessions and used by the dialogue engine to support relational behaviors (Tables 8-9 and 8-10 give examples of some of these persistent properties).
Alerts	Events the FitTrack system operator should be aware of and take action on, including system anomalies, emotionally distraught subjects, and user feedback.

Table 8-1. FitTrack Database Schema Overview

When the client application is launched the Executive first attempts to make a session connection to the Dialogue Server, which is then maintained throughout the session and is disconnected when the user exits the application. Messages between the Dialogue Server and Executive are ASCII XML statements of the form specified in Tables 8-2 and 8-3. Following a successful connection, the client sends a USER_LOGIN message to the server and the server responds with a SESSION message. After this, the server sends PERFORM messages to control the agent and browsers and the client sends PERFORM_COMPLETE messages upon completion of each PERFORM, USER_INPUT messages following user input actions, and WEB_PAGE messages whenever the contents of the main browser window are changed.

XML Message	Description
<USER_LOGIN ID= <i>ID</i> PWD= <i>PWD</i> >	Sent when client is started, with user ID and password.
<PERFORM_COMPLETE/>	Sent when PERFORM script has finished execution
<WEB_PAGE URL= <i>URL</i> />	Sent when browser changes web page.
<USER_INPUT TEXT= <i>text</i> />	Sent when user inputs text into text entry box.
<USER_INPUT MENU= <i>item</i> />	Sent when user selects a menu item.
<USER_EXIT/>	Sent when client application terminates.
<TEST DAY=.. COND=.. USERID=.. />	Used to reconfigure server (for testing purposes only).

Table 8-2. Messages from FitTrack Client to Server

XML Message	Description
<code><SESSION ID=<i>id</i> /></code>	Sent following successful startup and validation.
<code><SESSION ERROR=<i>msg</i> /></code>	Sent following unsuccessful startup.
<code><FLUSH/></code>	Causes client to halt script execution and reset all buffers.
<code><PERFORM> <i>Command</i> * </PERFORM></code>	Instructs client to begin execution of new script.
<code><i>Command</i> ::= <i>InterfaceCommand</i> <i>AgentCommand</i></code>	Agent Commands are specified in Table 8-4.
<code><i>InterfaceCommand</i> ::= <i>DelayCommand</i> <i>CameraCommand</i> <i>PageCommand</i> <i>MenuCommand</i> <i>InputCommand</i> <i>SidebarCommand</i></code>	Commands a change to the client interface.
<code><i>DelayCommand</i> ::= <DELAY MS=<i>ms</i> /></code>	Delays script execution for the specified number of milliseconds.
<code><i>CameraCommand</i> ::= <CAMERA ZOOM=0..1/></code>	Commands camera to zoom in or out on the agent (zero is a full body shot. one is a close up on the agent's face).
<code><i>PageCommand</i> ::= <PAGE URL=<i>URL</i> /></code>	Commands the main browser window to load a web page.
<code><i>MenuCommand</i> ::= <MENU> <i>MenuItem</i>* </MENU></code>	Displays a multiple-choice menu with the specified items.
<code><i>MenuItem</i> ::= <ITEM> <i>text</i> </ITEM></code>	Specifies one item on a multiple-choice menu.
<code><i>InputCommand</i> ::= <INPUT PROMPT=<i>text</i> /></code>	Displays a text-input box with the specified prompt.
<code><i>SidebarCommand</i> ::= <SIDEBAR URL={<i>URL</i> NONE} /></code>	Displays the specified page in the sidebar (help) browser window, or hides the sidebar browser window (if URL is "NONE") to expose ECA.

Table 8-3. Messages from FitTrack Server to Client

8.3 FitTrack ECA Requirements and Specifications

In order for the animation system to work in real time on a standard home computer, a 2D vector-based rendering method was used to draw each frame of the animation. Animation segments were pre-compiled, for either the character's entire body or just selected body parts, then played in synchrony with a speech synthesizer at run time. The final design uses Macromedia Flash for rendering, driven by the Microsoft Whistler speech synthesizer in real time using event callbacks. Both of these tools run on typical PCs with minimal processing and memory resources, are very stable and robust, and the use of the Microsoft text-to-speech synthesizer enables speech to be delivered with very little latency between turns (as compared

to using Festival pre-utterance timings, for example, as in (Cassell, Vilhjálmsón, & Bickmore, 2001)). Full body and body part animations were first crafted in Creative Lab's Poser, then exported with the aid of custom scripts to Flash, where final hand-alignment and cleanup was performed.

The nonverbal behavior supported by the FitTrack ECA includes the following:

1. Two full-body postures for all speaking behaviors, with shifts between them to signal topic shifts (Cassell, Nakano, Bickmore, Sidner, & Rich, 2001).
2. Ability to display the character in any camera framing from a wide shot (full body) to a closeup (face only), to support perceived proximity manipulations. Flash supports this functionality directly given that it uses vector-based rendering.
3. Ability to walk on and off of the screen (to/from one posture, with no other ongoing behavior) to signal natural transitions into and out of conversation with the character.
4. Several short, very subtle, full-body and facial idle behaviors (self-adaptors, e.g., scratching, looking around, looking at floor) and eye blinking in all facial expressions (brows down, closed mouth) to give the character more natural behavior when it wasn't speaking.
5. Conversational planning look-away and return (within each posture, with no other ongoing behavior), to signal turn-taking and turn-holding.
6. Four facial expressions: neutral, warm, concerned and happy. Three of these--neutral, concerned and happy--are from observations of the fitness trainer in Pilot 1 (chapter 7), with two versions of a neutral face implemented for RELATIONAL (warm face) and NON-RELATIONAL (neutral face) conditions. While the character can talk and raise and lower eyebrows in any facial expression, transitions between expressions are only supported in the brows down, closed mouth state. Thus facial expressions can only be changed at clause boundaries.
7. Head nodding in all facial expressions (brows down, closed mouth) to signal emphasis and acknowledgement.
8. A standard set of "visemes" (closed, A, B, E, F, K, O, OO, R, TH) in all facial expressions for natural lip-synchronization.
9. Eyebrows raise and lower in all facial expressions (during any viseme) to indicate emphasis.
10. A range of hand gestures that can be used during and synchronized with speech, including co-articulated beat gestures (i.e., beat gestures while in the middle of deictic or contrast gestures) (see Table8-4).

To support these functions with the minimum number of animation clips, rendering is seamlessly switched between the use of full body animations (used when synchronization among all body parts is of paramount importance) and composited animations comprised of head and arm animation clips that can be individually manipulated. The animations are layered as shown as in Figure 8-3. The Executive keeps track of the state of each body part, and when it receives a new animation command it moves displayed parts into the required states in which transitions from full-body to composited views are possible, when necessary, to support the animation command. Even with this compositing, however, approximately 250 separately-addressable animation clips are required to support the functions specified above.

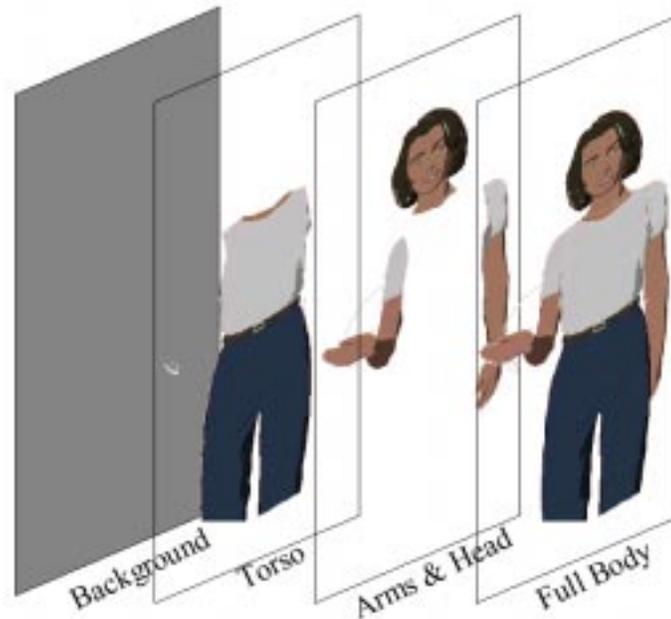


Figure 8-3. Animation Layering for FitTrack ECA

Table 8-4 specifies the command language used to control the ECA. Coverbal behavior synchronization is specified by embedding nonverbal behaviors in the utterance at the location they are to occur during production. For example, the following command causes the agent to perform a beat gesture when the word "walking" is spoken:

```
<SPEECH> Try <GESTURE HAND=L CMD=BEAT> walking.
</GESTURE></SPEECH>
```

8.4 Exercise Advisor Agent "Laura"

The appearance and nonverbal behavior of the exercise advisor was based on a review of relevant literature and a series of pre-test surveys.

In studies of the effects of appearance and nonverbal behavior in task settings it has been found that people are more responsive to those who are similar in appearance, and react more favorably (e.g., are more helpful) to those who are respectably dressed and physically attractive (Argyle, 1988). In addition, there have been many studies of the effects of the appearance and nonverbal behavior of psychotherapists on therapy outcomes (Beutler, Crago, & Arizmendi, 1986). One of the most relevant findings are that, in general, age-similarity is preferred, but especially among patients of a younger age, therapists of similar age to their clients have better outcomes. With respect to therapist gender effects on outcomes, it was found that female therapists, first, and gender-matched therapists, second, facilitated a greater degree of "treatment benefit". There have been no substantiated effects of therapist ethnicity on outcome. As previously mentioned, immediacy behaviors have been found to positively influence outcomes in both psychotherapist-patient and physician-patient interactions (Beutler et al., 1986; Buller & Street, 1992).

Command	Description
<i>AgentCommand ::= NVB Speech</i>	
<i>Speech ::= <SPEECH> { word SNVB }* </SPEECH></i>	Commands agent to speak.
<i>NVB ::= PostureCommand GazeCommand FacingCommand HeadnodCommand FaceCommand DisplayCommand AnimateCommand SNVB</i>	Commands agent (or browser) to execute a non-speech action.
<i>PostureCommand ::= <POSTURE/></i>	Commands agent to perform a posture shift.
<i>GazeCommand ::= <GAZE DIR={AWAY TOWARDS}/></i>	Commands agent to gaze away from or towards the user.
<i>FacingCommand ::= <FACING DIR={AWAY TOWARDS}/></i>	Commands agent to face away from or towards the user.
<i>HeadnodCommand ::= <HEADNOD/></i>	Commands agent to perform a headnod.
<i>FaceCommand ::= <FACE EXPR=expr /></i>	Commands agent to change its facial expression. Implemented expressions include NEUTRAL, WARM, SMILE, and CONCERN.
<i>DisplayCommand ::= <DISPLAY CMD={SHOW HIDE}/></i>	Commands agent to walk on or off screen.
<i>AnimateCommand ::= <ANIMATE ID=id /></i>	Commands agent to perform a specified animation sequence.
<i>SNVB ::= EyebrowCommand GestureCommand</i>	Specifies an agent nonverbal behavior that can be executed synchronously with speech.
<i>EyebrowCommand ::= <EYEBROWS DIR={UP DOWN}/></i>	Commands the agent to raise or lower its eyebrows.
<i>GestureCommand ::= <GESTURE HAND={L R} CMD=GestureType /></i>	Commands the agent to perform a hand gesture with the specified (Left or Right) hand. GestureTypes include: RELAX -- Hand at side. READY -- Hand in gesture space. BEAT -- Beat gesture (from any position). CONTRAST -- Contrast gesture. DOWN -- Deictic down. LEFT -- Deictic left. YOU -- Deictic at user. ME -- Deictic at agent (self). OK -- Emblematic "thumbs up" gesture.

Table 8-4. Exercise Advisor PERFORM Commands to Control the ECA



Figure 8-4. Candidate Exercise Advisor

These findings indicate that the exercise advisor agent should be female, approximately the same age as a typical college student (expected to make up the bulk of our user base), physically attractive, and respectably dressed. Based on these criteria a professional animator designed approximately one dozen candidate characters in the Poser 3D modeling software package. These candidates were then reduced to a set of three finalists through discussions with other students and thesis committee members (see Figure 8-4).

8.4.1 Character Selection Survey

A web-based survey was conducted within the Media Lab to select a final look for the exercise advisor character. Subjects were recruited via the Media Lab internal mailing list, and 41 responses were recorded. Subjects were shown the 3 characters shown in Figure 8-4, and asked to rank order them based on how comfortable they would be working with the character on a daily basis. Subjects were also asked about their stage of behavior change with respect to the ACSM/CDC standards for minimum activity level in order to determine the opinions of study candidates (17 of the 41 respondents were classified in contemplation, preparation or action and were thus "candidates"). Results from this study are given in Table 8-5. The characters (Advisors 1, 2 and 3) were rated on a scale of "Most Comfortable" (rating = 1), "Second Most Comfortable" (rating = 2), and "Least Comfortable" (rating = 3).

Given that the purpose of this survey was primarily to eliminate characters that some subjects might feel uncomfortable with (e.g., because the character is too thin, or mis-matched in age, ethnicity or gender), there was a clear preference among all respondents for Advisor 2, while candidates rated Advisors 1 and 2 approximately equal with respect to comfort. Consequently, Advisor 2 was selected for further development.

	Advisor 1	Advisor 2	Advisor 3
Respondents average rating	2.17	1.71	2.12
Respondents who rated as "most comfortable"	12	19	10
Respondents who rated as "least comfortable"	19	7	15
Candidates average rating	1.88	1.94	2.18
Candidates who rated as "most comfortable"	7	5	5
Candidates who rated as "least comfortable"	5	4	8

Table 8-5. Results from Character Survey

8.4.2 Nonverbal Behavior Survey

A second web-based survey was then prepared to rate a variety of body postures and facial expressions on how "warm" or "cold" they were perceived to be. Although the character will be in fairly constant motion during the interaction with the user, these postures and facial expressions were intended to be used as baselines onto which other animations would be layered. The intention was to select a small number of postures and facial expressions from those rated most "warm" for use in the RELATIONAL condition of the study, and a small number of postures and facial expressions from those rated "neutral" (in between warm and cold) for use in the NON-RELATIONAL condition of the study.

Twenty-two subjects responded to the survey, providing ratings and comments on 7 facial expressions and 9 body postures. Most of the body postures were presented twice to subjects (though never consecutively); once with a cold/neutral face and once with a warm face. Ratings could range from 1 ("cold") to 4 ("neutral") to 7 ("warm"). Results are presented in Table 8-6.

In addition to providing rankings on each of the individual postures and facial expressions, several interesting findings emerged from analysis of the data:

- The effects of consistency between head and body were as expected: a warm face ("Smirk") with a warm body ("Relaxed") received the overall warmest rating (5.59); a cold face ("Neutral") with a cold body "Hands-On-Hips") received the overall coldest rating (2.91); the inconsistent cases fell in between with the face carrying more weight (warm head/cold body = 4.59; cold head/warm body = 3.64).
- The face seemed to play a strong role in perception of the body. Scores averaged over postures with the warm face were significantly warmer (5.19) than scores averaged over postures with the neutral face (3.61). In every case in which a posture was rated with a warm and cold face, the body with the warmer face was rated as warmer.
- Overall, close up shots (Near proximity) are perceived as more warm (4.61) than body shots (4.08; a "proximity effect").

Based on the rankings provided in this study, the three facial expressions observed in Pilot Study 1 (described in chapter 7), and input from an expert on facial expressions (Nancy Alvarado, IBM Research), the following postures and facial expressions for selected for use in the final system:

- The "Neutral" and "Relaxed" postures (see Figure 8-5) are used in both RELATIONAL and NON-RELATIONAL conditions, since these can be perceived as either warm (5.59, 5.73) or neutral (2.50, 3.64) depending on the facial expression they are used with. Only two postures are needed to shift between in order to indicate discourse segment shifts.

Proximity	Posture	Face	Mean	StdDev
Far	Hands-on-Hips	Neutral	2.91	1.27
Far	Hands-on-Hips	Smirk	4.59	1.22
Far	Meek	Neutral	4.14	1.21
Far	Meek-Rotated	Neutral	2.77	1.07
Far	Meek	Smirk	4.86	0.77
Far	Neutral	Neutral	2.50	0.96
Far	Neutral	Smirk	5.73	1.12
Far	Relaxed	Neutral	3.64	1.18
Far	Relaxed	Smirk	5.59	0.73
Near		Concerned	5.04	1.22
Near		Neutral	3.52	1.12
Near		Neutral/Straight	3.87	0.69
Near		Smile	5.87	0.69
Near		Smile/Straight	5.50	0.74
Near		Smirk	5.96	1.26
Near		Superior	2.48	1.27

Table 8-6. Results from Nonverbal Warmth Survey

- The baseline facial expression in both the RELATIONAL and NON-RELATIONAL conditions is the "Smile/Straight" face with a minor modification.¹ Although the "Smirk" face was rated as the overall warmest, the facial expression expert was concerned that the expression could be mistakenly read as sarcastic or judgmental (e.g., in situations in which the advisor is giving evaluative feedback), exacerbated by the extreme head tilt. The "Smile/Straight" face was modified by slightly raising the eyebrows and widening the eyes (per the expert's suggestion) to convey more of a sense of warmth.



Figure 8-5. Two body postures used in the final system .

¹ Just before the study presented in chapter 9 was started, a concern was raised that the difference between a baseline neutral and baseline warm face may convey more than just relational stance. Subsequently, the same, warm face was used for both RELATIONAL and NON-RELATIONAL conditions.



Figure 8-6. Exercise Advisor Facial Expressions. Clockwise from upper left: Neutral; Warm; Happy; Concerned

- The "Concerned" face is used in both RELATIONAL and NON-RELATIONAL conditions when providing empathetic feedback to the user.
- The "Smile" face is used in both RELATIONAL and NON-RELATIONAL conditions when expressing pleasure/happiness or showing signs of increased immediacy (e.g., during social dialogue).

Figure 8-6 shows the final four facial expressions implemented for the exercise advisor (note that the "Neutral" face was not used in the study presented in chapter 9).

8.4.3 Facial Expression Survey

Another survey was conducted to ensure that the facial expressions selected for use were perceived to carry the emotional expression intended. Subjects were shown each of the four images shown in Figure 8-6 and asked to rate them on four seven-point semantic differential scales: Engaged/Disengaged, Warm/Cold, Concerned/Unconcerned, and Happy/Sad. Fourteen subjects responded; results are shown in Table 8-7.

Face	Engaged	Warm	Happy	Concerned
NEUTRAL	4.46	4.00	3.38	4.38
WARM	5.62	5.92	6.15	3.54
HAPPY	5.85	5.54	6.54	3.46
CONCERNED	5.58	5.08	2.92	6.25

Table 8-7 Ratings of Facial Expressions

8.4.4 Character Name Survey

A final survey was conducted to find a suitable name for the exercise advisor character. The primary concern was to find a name that had low probability of carrying specific connotations for users. The search started with a list of the 20 most common female baby names from 1980 (the approximate time that a 20-25 year old person would have been born) and, through several discussions with colleagues, this list was narrowed down to six candidates: "Melissa", "Tracy", "Laura", "Mary", "Julie" and "Karen". Another web-based survey was then conducted in which subjects were shown a picture of the character (Warm face, Neutral body, full body shot) and asked "Which of the following do you feel is the most appropriate name for this character?" Results indicated that "Laura" and "Karen" tied for the lead (each had 7 of 26 subjects voting for it). An executive decision was then made to use the name "Laura".

8.5 Exercise Advisor Server

As shown in Figure 8-2, the server side of the Exercise Advisor system is divided into two parts--the web server and the dialogue server--which communicate with each other via the relational database.

8.5.1 Web Server

The web server is primarily responsible for providing the correct sequence of pages (activity log forms, questionnaire forms, self-monitoring plots, and educational content pages) to a client browser for each user on each day and storing the results of any forms filled out by users in the database. The web forms and sample plot and educational content pages are given in Appendix B. The sequencing of pages is controlled by a finite state machine that goes through the following major steps each time a user logs in:

1. Activity logs are presented for the user to fill out, including hours of sleep, estimated minutes of physical activity in three intensity categories (following (Sallis, 1997)), and pedometer reading. These begin seven days prior to the user's first log in day (for an initial seven day recall to establish a shaping baseline) and continue, in order, through the day before the current day. If a user doesn't log in for a day or two, the system presents activity logs for the missed days. If the current time is after 2pm, the user is asked if they expect to do any more exercise today and, if they indicate that they do not expect to, an activity log for the current day is presented.
2. Self-monitoring charts are presented. Two separate charts and legends are presented for minutes of moderate-or-better activity and pedometer step counts (see Figure 9-4). The legends include: the user's "Ultimate Goal", which is the activity level they are working up to (typically 30 minutes or 10,000 steps per day); the user's "Current Goal", which is their current location on the shaping curve (a simple linear increase from baseline to ultimate goal from weeks 2-3); the user's "Current Commitment", which is how much they told the exercise advisor they would do during their previous interaction; and "Best Since Last Login", which is the best daily performance the user has had since the last time they logged in.
3. Any questionnaire forms that need to be filled out are then presented to the user. These are indexed by the user's interaction day (the number of calendar days since the user first logged in), and include any questionnaires required by the study (described in

- section 8.3). If the user misses a day on which a questionnaire was supposed to be administered, the system presents the questionnaire to them the next time they log in.
4. For users in RELATIONAL and NON-RELATIONAL treatments, the left-hand browser window then transitions into a Flash display, the exercise advisor agent walks on and begins a brief conversation with them. When the conversation is complete, the agent walks off and the left-hand browser window re-appears. Conversations are indexed by interaction number, rather than days since logged on, so that a logical sequence of conversations can be performed. One implication of this is that most users do not engage in all 30 conversations; each user only engages in 30 minus the number of days they were not able to log in.
 5. An educational content page is then displayed, indexed by interaction number. These pages--developed in conjunction with the MIT Medical Health Education Department--are designed to provide 1-5 paragraphs of information about an exercise-related-topic, sequenced to support someone beginning a new exercise program. Material was drawn from several governmental and nonprofit organization web sites, including pages from the NIH, CDC, and ACE (the American Council on Exercise). Table 8-7 shows the list of content pages developed.
 6. Once a content page is displayed, a user can spend as much time as they like reading the page. A hypertext link is also provided to an index of all content pages, any of which can be viewed at any time. Links are also provided that enable a user to re-display their self-monitoring chart or a feedback form, which can be used at any time to report problems with the system or communicate with the operator.

The web server logic is also responsible for determining when a user has finished the intervention part of the study (a page is displayed telling subjects that they are finished with the first part of the study, and to wait until they receive an email before logging in again) or are ready for the final follow up forms (activity logs and questionnaires).

Although users are told to log in once per day, the system is designed to support any number of logins per day, at any time of day, for any duration of time. Similarly, even though subjects are told they can only miss up to three consecutive days without logging in, the system is designed to support absences of arbitrary lengths of time and still present all activity logs and questionnaires required to make a user's data current the next time they log in. Once a user has logged on once in a given day, on subsequent logins they will typically only see their self-monitoring chart (step 2) followed by the educational content page (steps 5 and 6). In particular, users in RELATIONAL and NON-RELATIONAL conditions can only have one conversation with the agent on any given day.

8.5.2 Dialogue Server

The dialogue server maintains a persistent connection with a client application as long as the client is running, and records all client actions in the database for subsequent analysis. When a client first connects to the dialogue server, the client sends the user's ID and client application version number to the server. Following successful validation of the ID, the server creates a Session record in the database and instructs the client's browser to display the log-in form, which allows users to enter their password. Once this form is submitted to and validated by the web server, the Session record is updated to indicate that the user's session is active

Interaction	Topic
1.	Why Walk?
2.	What makes a walk a workout?
3.	Physical Activity Recommendation
4.	Pedometers
5.	How to Start and Maintain a New Exercise Program.
6.	Warm-up and Conditioning Exercises
7.	Health Benefits of Walking
8.	Staying Hydrated
9.	Safety
10.	Overcoming Excuses for not Exercising
11.	The Right Shoes
12.	The Right Clothes
13.	Staying Motivated
14.	The Best Time of Day to Exercise
15.	Target Heart Rate
16.	How to Get Energy for Exercise
17.	Correct Posture: The Stride
18.	Correct Posture: The Arms
19.	Preventing Injuries: Ankles
20.	Preventing Injuries: Shin Splints
21.	Burning Calories: Short Term
22.	Preventing Injuries: Feet
23.	Burning Calories: Long Term
24.	How to Walk Faster: Lower Body
25.	How to Walk Faster: Upper Body
26.	Breathing and Walking
27.	Walking to Music
28.	More Stretches to Try
29.	Walking Away Stress
30.	Treadmills

Table 8-7. Educational Content Page Topics and Sequencing

(until this is done, no other parts of the system are functional). At this point, the web server has primary control, and delivers web forms to the user as described above. During this time, the dialogue server is simply monitoring what is going on and recording client interface events in the database. However, once a particular web page is loaded (telling the subject they are going to have a discussion with the agent) the dialogue server initializes the dialogue state machine for the current interaction, causes the agent to walk on screen and begins the conversation with the user. Once the conversation is complete, the dialogue server instructs the agent to walk off the screen, and resumes monitoring client interface actions until the user exits the program.

The dialogue server is driven by another state machine that keeps track of high-level dialogue events, such as whether the character just asked the user a question and is waiting for a response, or whether the user just quit the client application. A significant portion of this

state machine is designed to handle error conditions, including the case when a subject stops interacting with the agent in the middle of a conversation. In this case, the agent repeats its last utterance after a 30 second timeout, if the user still does not respond after two minutes the agent asks "Are you there?", if there is still no response after another minute, the conversation is ended and the agent walks off the screen. If any internal errors are encountered during a conversation, the errors are logged in the database, the agent says "goodbye" and the conversation is ended. Figure 8-7 shows a simplified view of this state machine in which all of the error handling states have been removed. Discussion of this state machine is covered in the next section.

8.6 Dialogue Scripting Language and Compiler

Designing and implementing 30 ECA conversations is a significant undertaking, even when the conversations only last 5-10 minutes each. In the Pilot 1 conversations presented in Table 7-3, turn lengths averaged 2.8 seconds each, indicating that approximately 80 agent utterances will need to be implemented for each conversation and over 2,000 utterances implemented for the 30 days, and this is assuming that users are only given the option of saying one thing at any point in the conversation. Allowing for a range of user responses and providing enough system variability to make the agent seem lifelike increases the complexity of this task even further.

To support a dialogue system of this magnitude and complexity, an Augmented Transition Network (ATN) approach was selected for its design. Augmented Transition Networks were designed for use in parsing natural language utterances (Woods, 1986), and involve the use of a collection of finite state machines and a single runtime stack. Execution begins at a start node in one of the state machines, but in addition to branching to another state in the same machine, execution can branch to the start node of another state machine by pushing the current state onto the stack. When the end state of the second network is reached, the stack is popped and execution resumes in the first network at the state that was popped from the stack. This mechanism allows common dialogue fragments to be "factored out" into separate finite state machines and re-used in several conversations to reduce the overall complexity of the network.

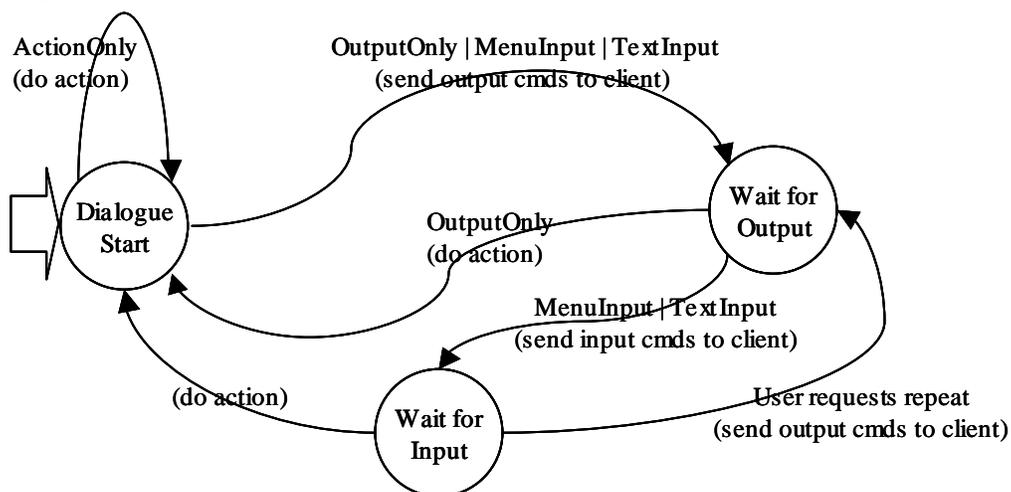


Figure 8-7. Fragment of ATN Executive State Machine

To reduce the complexity of implementing a large set of dialogue state machines, a scripting language was designed to allow them to be described as concisely as possible, and then compiled directly into executable code as part of the Dialogue Server. The syntax for the scripting language is shown in Table 8-8. An example state described in this language is shown below. In this example (a specification for state "GC_16"), when the state is entered the agent will ask the user "Are you going to work out tomorrow?", and the user is presented with the menu choices "I can't.", "Yep." or "Could you repeat that?" (an option requesting that the agent repeat its last utterance is automatically added to the end of each menu). Execution then branches to state "GC_19" or state "GC_18" depending on which menu item the user selects.

```
STATE: GC_16
AGENT: $ <BEAT>Are you going to work out tomorrow? </BEAT> $
USERMENU: I can't. => $ GO("GC_19"); $
          Yep. => $ GO("GC_18"); $
```

In addition to 'GO' actions, actions in the scripting language can include the following primitives:

- PUSH(FSA2Start,FSA1Next) -- Pushes state 'FSA1Next' from the current finite state machine onto the runtime stack, then begins execution of state 'FSA2Start' (the start state of finite state machine 2).
- POP() -- Pops the runtime stack and continues execution at the popped state.

Referring back to Figure 8-7, the start of execution for each dialogue state begins in the "DialogueStart" state. If the script associated with the state specifies any output commands, these are sent to the client and the machine transitions to the "WaitForOutput" state. In the example above, these output commands would consist of a PERFORM specifying SPEECH plus nonverbal behaviors added in by BEAT, which would look something like the following:

```
<PERFORM><SPEECH>Are you going to work out tomorrow?
</SPEECH></PERFORM>
```

Once the client has responded with a "<PERFORM_COMPLETE/>" message (upon completing execution of the output commands), the input commands to configure the client's input menu or text dialogue box are sent from the dialogue server to the client, and the machine transitions to the "WaitForInput" state (assuming a MenuInput or TextInput script node). In the example above, the input commands would be a message of the following form:

```
<PERFORM><MENU><ITEM>I can't.</ITEM><ITEM>Yep.</ITEM>
<ITEM>Could you repeat that?</ITEM></MENU></PERFORM>
```

The machine waits in this state for the user to input either a menu choice or text. Once the client sends this data (via a "<USER_INPUT MENU=.../>" message) the machine executes any actions associated with the script, conditioned on the user's input. These actions typically include changes to the dialogue state (either through an ATN branch, push or pop operation), after which the machine transitions back to the "DialogueStart" state. In the example above, if the user selected "Yep." from the input menu, a message of the form "<USER_INPUT

Syntax	Description
<code><network> ::= <state>*</code>	The states for a single finite-state machine network are defined in a single file, with equivalent labels for the network, first state and filename.
<code><state> ::= 'STATE:' <stateLabel> { <actionOnlySpec> <outputOnlySpec> <menuInputSpec> <textInputSpec> }</code>	Each state is one of four types: 1) action only; 2) output only; 3) menu input; or 4) text input.
<code><actionOnlySpec> ::= <action></code>	In an action-only state, the server performs an internal action then branches to a new state.
<code><outputOnlySpec> ::= <agentOutput> <action></code>	In an output-only state, server sends commands to client and branches to a new state.
<code><menuInputSpec> ::= <agentOutput> <menuInput></code>	In a menu-input state, the server sends commands to client, including a multiple-choice menu specification, waits for user input, then branches to a new state.
<code><textInputSpec> ::= <agentOutput> <textInput></code>	In a text-input state, the server sends commands to client, including a text input specification, waits for user input, then branches to a new state.
<code><action> ::= '\$' <java> '\$'</code>	An action consists of java language code, including a number of primitives implemented specifically for the dialogue server.
<code><agentOutput> ::= { <agentBoth> <agentRel> <agentNRel> }+ [<agentRepeat>]</code>	Commands to send to the client can be specified for all users ('AGENT:') only users in RELATIONAL condition ('AGENT_REL:') or only users in NON-RELATIONAL condition ('AGENT_NREL:'). In addition, a 'REPEAT:' clause can be specified to give variations of output if the user asks the agent to repeat the last utterance. Note that multiple outputs can be specified for any or all of these categories, in which case the server picks one randomly at runtime to give variety in the agent's responses.
<code><agentBoth> ::= 'AGENT:' { '\$' <agentCmds> '\$' }+</code>	
<code><agentRel> ::= 'AGENT_REL:' { '\$' <agentCmds> '\$' }+</code>	
<code><agentNRel> ::= 'AGENT_NREL:' { '\$' <agentCmds> '\$' }+</code>	
<code><agentRepeat> ::= 'REPEAT:' '\$' <agentCmds> '\$'</code>	
<code><menuInput> ::= 'USERMENU:' <userMenuResponse>+</code>	Specifies a multiple-choice selection menu and the actions to take when an entry is selected. The server always adds an item to the end of each menu that allows the user to request that the agent repeat its last utterance.
<code><userMenuResponse> ::= <userMenuWords> { ';' <userMenuWords> }* '=' <action></code>	
<code><textInput> ::= 'USERTEXT:' <userPromptWords> '=' <action></code>	Specifies that a text input box and prompt be displayed, and the actions to take when the user enters something.
<code><agentCmds> ::= XML</code>	Specifies XML commands to send to the client (defined in Tables 8-3 and 8-4). In addition, any text put between '<BEAT>' tags is pre-compiled through BEAT, and '=' <javaExpr> '=' is evaluated in-place at runtime with the result spliced in before being sent to the client.

Table 8-8. Dialogue Server Scripting Language Specification

MENU=2/>” is sent from the client to the dialogue server, indicating that the second menu option had been chosen. The ATN executive would then execute the action ‘GO(“GC_18”)’ causing a branch to the ATN dialogue state labeled ‘GC_18’. At this point the process would start over with the ATN executive in the “DialogueStart” state and the current ATN node set to ‘GC_18’. Once the ATN stack is empty, the dialogue ends.

A compiler was developed that takes the dialogue scripts as input and produces executable code as output. The compiler first parses the script files and builds an intermediate representation. Next, it performs a number of checks on the state machine specification to ensure there are no cycles, unreachable or non-existent states, or common errors such as specifying output for RELATIONAL but not NON-RELATIONAL conditions. The agent output utterances are then pre-processed through a modified version of BEAT (described in chapter 6) which converts BEAT input strings (utterances with nonverbal behavior markup tags) into PERFORM primitives (described in Tables 8-3 and 8-4) that can be sent directly to a client for execution. The BEAT tags for relational stance are derived from the user’s experimental condition specified in the database (i.e., RELATIONAL users have STANCE set to “WARM”, NON-RELATIONAL users have STANCE set to “NEUTRAL”). BEAT tags for conversational frame (“SOCIAL”, “EMPATHY”, etc.) are specified manually in the scripting language.

One additional function built into the scripting language is the ability to perform computations at runtime and splice the results into output utterances just before they are sent to the client. This is used, for example, to insert the user's name into a greeting utterance or look up information about them in the database and use it to tailor an utterance (e.g., "Did you play TENNIS again today?").

8.7 Persistence

An important part of relationships is being able to talk about past interactions and history together and to demonstrate “common ground” by remembering details of one’s relational partner. In order to support this, and to provide a mechanism for communicating information between dialogue state machines, a set of <user, property, value> tuples is maintained for each user in the dialogue server. These tuples are available for reading and writing from all dialogue states (thus supporting inter-state communication), and are saved in the database at the end of each dialogue and re-loaded from the database at the start of each dialogue (thus supporting recall of information discussed in prior conversations). In addition, a number of tuples are computed at the start of every dialogue to support common exercise advisor dialogue tasks.

There are 86 properties used in the final dialogue engine; Table 8-9 lists the properties that are initialized at the start of each session and Table 8-10 lists the most-commonly used persistent, single-valued parameters, whose values are saved across sessions. In addition, for many properties it is useful to know not just a single value, but a history of all values indexed by conversation. For example, it is useful for the agent to know which stories (third person stories used to relate exercise tips) have been told and when they were told, so that it could preface a story with "Remember last Tuesday when we talked about your not having enough time to exercise?". To support this functionality, some properties are represented by tuples of the form <user, property, interaction-number, value> . A list of the historical properties used in the final dialogue engine is given in Table 8-11.

Property	Description
STUDY_DAY	Current day of the study.
INTERACTION	Dialogue number.
CONDITION	Experimental condition.
DAY_OF_WEEK	"Monday", etc.
PART_OF_DAY	"Morning" "Afternoon" or "Evening".
NAME	User's name.
REPORT_AGE	Whether the latest activity report is for yesterday or today.
EXERCISE_TIME	Actual time of moderate-or-better activity on most recent activity report.
EXERCISE_STEPS	Actual number of steps on most recent activity report.
EXERCISE_BEST_TIME	Best daily time of moderate-or-better activity since last login.
SLEEP	Minutes of sleep on most recent activity report.
MODERATE_TIME	Minutes of moderate activity on most recent activity report.
HARD_TIME	Minutes of hard activity on most recent activity report.
VERY_HARD_TIME	Minutes of very hard activity on most recent activity report.
GOAL_TIME	Current day goal for minutes of moderate-or-better activity.
GOAL_STEPS	Current day goal for steps.
NEXT_EXERCISE_DAY	"Today" (if user indicated they are going to exercise more today) or "Tomorrow".
ULTIMATE_GOAL	Daily minutes of moderate-or-better activity the user is trying to work up to.

Table 8-9. Properties Initialized at the Start of Each Session

Property	Description
EXERCISE_GOAL	The user's most important reason for exercising.
EXERCISE_BENEFIT1	The benefits that the user expects to get from exercising.
EXERCISE_BENEFIT2	
ISLONER	Whether the user prefers to work out alone or with others.
BUDDY	Name of workout buddy.
COMMITMENT_TIME	Commitment to exercise made by the user.
COMMITMENT_TYPE	
OBSTACLE	The user's current obstacle to meeting their exercise goals.
UPCOMING_EVENT	Information about any big events coming up for the user (e.g., a final exam or project at work) and when they are happening.
WEEKEND_EVENT	What the user plans to do for the weekend.
RESIDENCE	Where the user lives.
LIKESMOVIES	Topics of social dialogue.
LIKESTV	
CUISINE	
ORIGINS	
OCCUPATION	Information about the user.
STUDENT_YEAR	
STUDENT_MAJOR	

Table 8-10. Example of Some Persistent, Single-Valued Properties

Property	Description
STATE	Overall physical/emotional state of subject (e.g., "UPSET", "INJURED", "SICK").
FEELING	Subject's emotional state (e.g., "ANGRY", "FRUSTRATED", etc.).
ILLNESSTYPE	Information about the type of illness the subject has had.
DID_EXERCISETYPE	Information about what the subject did for exercise.
DID_EXERCISEDETAIL	
DID_EXERCISETIME	
DID_EXERCISEBUDDY	
FEELING_PERSON	If the subject's emotional state is due to another person, this is the name of that person.
STORIES_TOLD	List of exercise "stories" told to the subject.
COUNTERARGS_TOLD	List of obstacle counter arguments told to the subject.
RELAPSE	Record of which sessions the subject was in "relapse" (exercised below recommended goal level).
WEATHER	Record of the weather conditions, as reported by the user.

Table 8-11. Persistent, Interaction-Based (Historical) Properties

8.8 Exercise Advisor Dialogue Design

The design of the exercise advisor dialogues was based on the two pilot studies (described in chapter 7) and the literature on exercise behavior change, working alliance, and relational communication. The format of the first interaction borrows heavily from the first Pilot study, in which the advisor spends most of the time asking the user about their exercise goals and talking about what they can expect in future interactions with the agent. Subsequent interactions take the form of typical interactions in the second Pilot study, following the general template below (with significant daily variations):

1. Greeting
2. Determine the physical and emotional state of the user.
3. Social dialogue.
4. Follow up on the user's commitment to exercise made in the previous interaction.
5. Offer exercise tips.
6. Get the user to commit to a specific amount of exercise before the next interaction.
7. Farewell

The behavior change and relational literature each suggest separate schedules for phasing in different topics of conversation.

For behavior change, the agent spends early sessions discussing the user's reasons for beginning an exercise program and negotiating long-term exercise goals (what the user wants to achieve by the end of the program), as well as finding out information about the user's exercise preferences (e.g., whether they like exercising alone or with others, whether they play a sport, etc.). Other topics discussed in early sessions include important information such as stretching and how to walk to get maximum cardiovascular benefits, while later talks focus more on motivational issues. The last few interactions focus on maintenance and relapse prevention techniques. Throughout the intervention period, the agent works with the user to

negotiate daily goals based on their behavior shaping curves (the agent tries to ensure they don't commit to too little or too much relative to the curve).

As the working alliance is generally established within the first seven interactions between a helper and helpee, the relational strategies are all phased in by the seventh interaction. The amount and depth of social dialogue, humor, and meta-relational communication are gradually increased in interactions 1 through 7. References to future interactions, use of inclusive pronouns, expressing happiness to see the user, use of close forms of address (subject's name), politeness strategies, and empathy are used at a constant rate, whenever the opportunity for their use arises. In addition, just as people in face-to-face conversation always have the option of leaving the conversation at any point, users are always given the ability to "opt out" of social dialogue (e.g., with moves like "Actually, I'm in a little bit of a rush today.>").

The final dialogues for the 30 interactions consist of 76 finite state machines, with a total of 974 states (see Table 8-12). Samples of user's second interactions with the agent, in NON-RELATIONAL and RELATIONAL conditions, are shown in Figures 8-8 and 8-9, respectively.

8.9 Summary

The relational agent described in this chapter represents many firsts for an ECA-based interface. It is the first explicitly designed for multiple interactions with a user over a long period of time, exhibiting persistence, continuity, trajectories of change, and variability, applied both to task and to social-emotional content. It is the first to support full body nonverbal behavior synchronized with speech that is designed for use on low-end personal computers (Microsoft Agent and other commercial ECAs do not allow for synchronization of

1. A: Hi!
2. U: Hello.
3. A: So, now I'd like to ask you a little about your exercise behavior. Do you currently do any kind of sport?
4. U: No, I don't
5. A: Do you currently do any other kind of aerobic exercise regularly?
6. U: I go to the gym
7. A: And, do you think you would like to exercise by yourself, or with a buddy?
8. U: By myself
9. A: Great. So, did you manage to go for a walk since we last talked?
10. U: No, didn't get the chance.
11. A: OK.
12. U: Before we chat again, do you think you could get out and go for another walk? It doesn't matter how long or how far.
13. A: Sure.
14. A: Okay. Here is some information about how to make a walk a workout. Casual strolls don't quite count.
15. A: Remember, when we end our sessions, you can spend as much time as you like browsing this library. Just close the Fit Track application when you are finished.
16. A: Have a great morning, and the rest of your day too. Goodbye.

Figure 8-8. Sample NON-RELATIONAL Interaction (A=Agent; U=User)

1. A: Hi Chrissy.
2. U: Hi Laura, great to see you.
3. A: How are you doing today?
4. U: I'm not feeling so great.
5. A: I'm sorry to hear that.
6. U: Tell me more about how you are feeling
7. U: I'm sick
8. A: I'm sorry you're sick. What do you think is wrong?
9. U: I have the flu
10. A: You must be feeling miserable. Have you been to see a doctor?
11. U: Yes.
12. A: That's good. It can make you feel better just seeing a health professional about a problem. Do you think this illness might affect your ability to exercise?
13. U: Definitely.
14. A: I'm sorry to hear that. As soon as you feel well enough, you might try to get out for a short walk. The fresh air will do you good, and it will help you to stay in the habit.
15. A: So, are you a student?
16. U: Yes
17. A: Do you go to MIT?
18. U: No
19. A: Oh, what school do you go to?
20. U: Cambridge Continuing Education, right now. I am working toward a degree at Tufts.
21. A: Really? Are you a grad student, or an undergraduate?
22. U: Senior
23. A: Cool. What's your major?
24. U: English
25. A: So, now I'd like to ask you a little about your exercise behavior. Do you currently do any kind of sport?
26. U: No, I don't
27. A : Do you currently do any other kind of aerobic exercise regularly?
28. U: I do something else
29. A: What do you do?
30. U: I do tap dancing. I also walk to and from work.
31. A: And, do you think you would like to exercise by yourself, or with a buddy?
32. U: With a buddy
33. A : Great. So, did you manage to go for a walk since we last talked?
34. U: Yes!
35. A: That's wonderful. Before we chat again, do you think you could get out and go for another walk? It doesn't matter how long or how far.
36. U: Sure.
37. A: OK. Here is some information about how to make a walk a workout. Casual strolls don't quite count. Remember, when we end our sessions, you can spend as much time as you like browsing this library. Just close the Fit Track application when you are finished.
38. A: So, when can we get together again?
39. U: Tomorrow.
40. A: Great. Take care Chrissy , I'll talk to you tomorrow. Good night.

Figure 8-9. Sample RELATIONAL Interaction (A=Agent; U=User)

Network	States	Utterances	Inputs	Description
Int1 - Int30	438	238	264	Top-level dialogue structure for daily interactions (day 1-30)
AskBuddy	12	10	10	Talk about workout buddy.
AskEnjoyment	4	3	5	Ask if user is enjoying their exercise routine.
AskNewGoal	6	5	7	Ask if user would like to set a higher ultimate goal level.
ContinuingFeeling	5	3	3	Empathy dialogue if the emotional state has persisted for more than a day.
ContinuingFeelingDown	11	9	12	Empathy dialogue if the user has been feeling down for more than a day.
ContinuingIllness	7	5	2	Empathy dialogue if the user has been ill for more than a day.
ContinuingInjury	8	7	2	Empathy dialogue if the user has been injured for more than a day.
ContinuingStressed	6	6	2	Empathy dialogue if the user has been stressed out for more than a day.
ContinuingTired	5	3	0	Empathy dialogue if the user has been tired for more than a day.
DoAffirmation	4	8	0	Give self-affirmation (positive thoughts).
DoPerformative	1	9	0	Give performative (agent voicing confidence in user).
Farewell	9	19	4	
FindState	13	18	24	Determine the physical and emotional state of the user.
FollowUp	25	34	6	Determine if user has met their prior commitment to minutes per day of exercise.
FollowUpObstacles	5	4	9	Talk about any obstacles the user has had to exercise.
FollowUpSteps	4	4	0	Determine if user has met their prior commitment to steps per day.
GetCommitment	44	38	94	Negotiate a commitment for exercise.
GetExerciseInfo	14	14	32	Talk about what the user did for exercise.
GetGoals	27	28	39	Talk about the user's goals and reasons for wanting to exercise.
GetStepsCommitment	11	8	30	Negotiate a commitment for steps.
GiveTip	2	0	0	Give an exercise tip or story.
Greeting	5	15	4	
HandleAngryFeeling	11	10	14	Empathy dialogue if user is feeling angry.

Table 8-12. FitTrack Dialogue States

Network	States	Utterances	Inputs	Description
HandleDisappointedFeeling	11	10	14	Empathy dialogue if user is feeling disappointed.
HandleFrustratedFeeling	11	10	14	Empathy dialogue if user is feeling frustrated.
HandleUpsetFeeling	11	10	14	Empathy dialogue if user is feeling upset.
MotivateDuration	6	7	20	Motivate user to increase their duration of exercise.
MotivateToExercise	5	6	10	Motivate user to exercise.
MotivateToWalk	8	7	11	Motivate user to go for a walk.
NewFeeling	5	8	8	Empathy dialogue if user is in a negative emotional state.
NewFeelingDown	14	16	18	Empathy dialogue if user is feeling down.
NewIllness	11	9	10	Empathy dialogue if user is ill.
NewInjury	17	16	18	Empathy dialogue if user is injured.
NewPersonFeeling	19	21	35	Empathy dialogue if person is upset about something that someone did to them.
NewStressed	7	7	7	Empathy dialogue if user is stressed out.
NewTired	9	11	11	Empathy dialogue if user is tired.
ObstacleTalk	2	0	0	Talk about user's obstacles to exercise.
OfferCounterArgument	25	29	0	Positive thoughts to counter negative thoughts about exercise.
OngoingEvent	15	18	27	Talk about a significant event in the user's life.
PastEvent	13	14	23	Talk about a significant event in the user's life.
PastEventTalk	2	0	0	Talk about a significant event in the user's life.
PastWeekend	16	17	29	Talk about what the user did over the prior weekend.
RelapseTalk	8	7	8	Help user who has been consistently underperforming.
TellStory	52	56	0	Tell a third person story about overcoming obstacles to exercise.
UpcomingEvent	14	12	32	Talk about a significant event in the user's life.
UpcomingEventTalk	2	0	0	Talk about a significant event in the user's life.
UpcomingWeekend	14	19	18	Talk about the user is going to do over the next weekend.
Total	974	808	890	

Table 8-12. FitTrack Dialogue States (continued)

anything other than lip-movements while the agent is talking). It is the first designed to assist with health behavior change. And, it incorporates one of the largest dialogue models—in terms of utterances and topics—developed for an experimental dialogue-based interface. The next chapter presents an evaluation of this system in a field trial involving 100 subjects over a six week period of time.

Chapter 9

AN EVALUATION OF LONG-TERM RELATIONSHIPS WITH RELATIONAL AGENTS

Many experiments have been conducted to evaluate people's behavior interacting with animated interface agents, as well their attitudes towards these characters (Dehn & Mulken, 2000). The findings have mostly been that people find them engaging to work with, however, there have been no consistent effects found on task outcome measures, such as whether these interfaces actually help users perform tasks more efficiently or effectively relative to alternative interfaces. Further, the reason that people find character-based interfaces engaging may be simply due to a novelty effect, since these experiments are almost always performed in a single brief session with a user (although occasionally within-subject designs call for two or three sessions over consecutive days, e.g. (Smith, 2000)). No one has yet to determine the effects of interacting regularly with interface agents over long periods of time. Do users tire of them quickly? Are there things the agent can do to keep the user engaged? Can relational agents play a routine role in people's everyday lives and have real impact on their behavior?

This chapter presents a first-of-its-kind study in which users interact with an animated agent on a daily basis for a month. The experiment is intended to investigate whether the relational behaviors that people use to build and maintain long-term relationships can be used by relational agents to bring people into a trusting, caring relationship with them over a long period of time, and if this relationship, in turn, can be used to significantly effect a task outcome—in this case a change in the user's behavior.

9.1 Introduction

An evaluation of the exercise advisor agent ("Laura"; described in chapter 8) was conducted using an empirical experiment on human subjects. The study utilized a longitudinal, multivariate, two-factor, between-subjects design to evaluate the effects of different intervention strategies on the level of physical activity in free-living subjects over a six-week period of time. The study followed the standard pattern for a behavior change study with an initial baseline measurement of the behavior of interest (via a 7-day recall on the first day), followed by an intervention period (30 days), followed by removal of the intervention to check if and when the new behavior extinguished (14 days at which time a second 7-day recall was administered) (Sunde & Sandra, 1999) (see Figure 9-1). The target behavior in this study is the current ACSM/CDC recommended minimum level of physical activity: "Every US adult should accumulate 30-minutes or more of moderate-intensity physical activity on most, preferably all, days of the week" (Pate et al., 1995).

The study was designed to ensure that subjects interact with the system for a brief period of time every day to provide the agent with an opportunity to build and maintain a relationship with them. The study has three treatments: RELATIONAL, NON-RELATIONAL, and a baseline CONTROL condition. In all treatments subjects recorded their daily activity via self-report forms, using 7-day recall at the start of the experiment and the end of the follow up period, and daily recall during the balance of the first month. Subjects

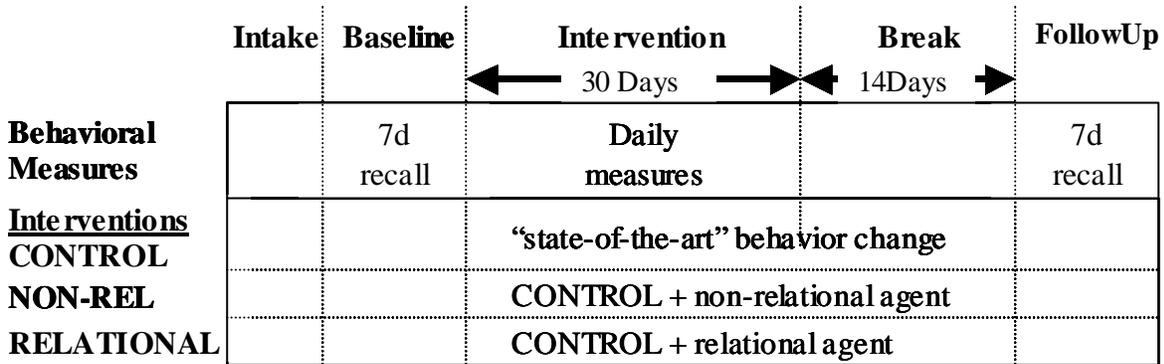


Figure 9-1. Overall Design of Study

were also given pedometers and reported the number of steps taken each day via a web form, to provide an objective measure of their physical activity level. In all conditions, subjects also received standard behavioral interventions, including self-monitoring (progress charts showing their activity levels over time) and decisional-balance questionnaires (e.g., as in (Nigg, Courneya, & Estabrooks, 1997)). All subjects were also provided with daily web pages of educational content on the topic of walking for exercise. All subjects in RELATIONAL and NON-RELATIONAL conditions also had a daily "conversation" with the virtual exercise advisor (Laura) about their progress, any obstacles they had to exercising, and the educational content, via a dynamic menu-input, synthetic speech output modality. In the RELATIONAL condition the agent also used relational strategies in an attempt to build a working alliance with subjects, whereas in the NON-RELATIONAL condition relational functionality has been ablated.

9.1.1 Hypotheses

The immediate effects of relational strategies by an agent should be an increase in measures of the quality of the relationship, such as the working alliance inventory.

H1. Subjects will feel more of a sense of relationship with the agent in RELATIONAL condition than in the NON-RELATIONAL condition.

As described in chapter 8, there is a significant literature indicating a strong correlation between working alliance and outcomes across a wide range of psychotherapeutic disciplines (e.g. (Horvath & Symonds, 1991)). Although no studies have been done relating working alliance to outcomes in exercise coaching, several studies have demonstrated an effect in other areas of health behavior change, including smoking cessation and alcoholism treatment. The working alliance has been demonstrated to have a significant correlation with outcome measures ranging from percentage of days abstinent, drinks per drinking day, and treatment participation (weeks in program) for alcoholism, to employment and compliance with medication, to more general measures such as premature termination, Global Rating Scale (GRS), Global Assessment Scale (GAS), MMPI, Cooley & Lajoy (C-L), and many, many others (Bachelor, 1991; Connors, Carroll, DiClemente, & Longabaugh, 1997; Gaston, 1990; Henry & Strupp, 1994; Horvath, 1994; Horvath & Luborsky, 1993; Horvath & Symonds, 1991; Keijsers, Schaap, & Hoogduin, 2000; Luborsky, 1994; Mallinckrodt, 1003; Raue &

Goldfried, 1994). One meta-analysis of 24 studies found an overall effect size of 0.26 on the relation between working alliance and outcome across a number of different psychotherapies (Horvath & Symonds, 1991). Based on these positive results, a significant effect of relational behavior on outcome is expected, as is the effect of the mere presence of an agent (increased sociality of the interface) on outcome.

H2 (a) Subjects in the RELATIONAL condition will perform more physical activity than those in the NON-RELATIONAL condition.

H2 (b) Subjects in conditions with the agent (NON-RELATIONAL and RELATIONAL) will perform more physical activity than those in the CONTROL condition.

Another measure of outcome that has been used in behavior change studies involving self-paced access to educational content is participation in the study, measured as frequency and duration of logins or number of educational pages viewed. Given that subjects in this study are able to log in as often as they like and spend as much time as they like viewing content pages, I would expect to see an association between working alliance (and the mere presence of a social agent) and measures of participation.

H3 (a) Subjects in the RELATIONAL condition will participate more in the study than those in the NON-RELATIONAL condition.

H3 (b) Subjects in conditions with the agent (NON-RELATIONAL and RELATIONAL) will participate more in the study than those in the CONTROL condition.

Finally, there is a strong correlation between certain attitudes towards a target behavior and the likelihood of the behavior being adopted. For example, self-efficacy (with respect to the target behavior) has been demonstrated repeatedly to be one of the single best determinants of ability to change (Dishman & Sallis, 1994). Thus, as increases in working alliance effect increases in the target behavior, I would also expect to see concomitant increases in attitudes towards that behavior.

H4 (a) Subjects in the RELATIONAL condition will have increased pro-exercise attitudes relative to subjects in the NON-RELATIONAL condition.

H4 (b) Subjects in conditions with the agent (NON-RELATIONAL and RELATIONAL) will have increased pro-exercise attitudes relative to subjects in the CONTROL condition.

9.2 Experimental Method

Table 9-1 gives an overview of the experimental protocol and measures used.

Day	Procedure
-	Subjects are recruited via email and flyers which direct them to a web page. Subjects are then screened via a web form for general level of health (ability to exercise), stage of change (contemplation, preparation and action (Marcus & Owen, 1992)), access to a personal computer with appropriate resources, and likelihood of completing all study tasks.
-	Subjects come into the laboratory for an intake meeting. They sign a consent form and are given initial instructions. Subjects are then given the screening questionnaire again, plus demographic and personality trait questionnaires. They are given a pedometer, written instructions, and CDROM with software to install on their home computers, which they are instructed to install as soon as possible.
1	Subjects fill out a 7-day recall of activity form via the web (the behavioral baseline measure).
1, 30, 44	Subjects fill out exercise self-efficacy and decisional balance via web forms.
1-30	Exercise adoption intervention. The system recommends 30 minutes of moderate-or-better activity most days of the week, with walking recommended as the exemplar. Walking 10,000 steps a day is also recommended as a secondary goal, since it roughly corresponds to 30 minutes a day of moderate activity (Tudor-Locke). Subjects also have access to self-monitoring charts and educational content pages, and fill out daily recall of activity forms. Subjects in RELATIONAL and NON-RELATIONAL conditions also have 3-5 minute conversations with Laura, the animated exercise advisor.
7-21	Shaping. Subjects' daily goal levels are increased from their baseline level (7-day recall for minutes activity and first three pedometer measurements for steps) up to the goal level, linearly.
7, 27	Subjects in all conditions except CONTROL fill out the Working Alliance Inventory form.
8	Subjects who have maintained more than 30 minutes of moderate-or-better activity in both their baseline week and first week of intervention are given the option to raise their goal level from 30 minutes per day up to 45, 60, 75, or 90 minutes per day (via web form).
31-43	Relapse check interval - no required activity. Subjects in RELATIONAL and NON-RELATIONAL conditions who give their permission are contacted for a qualitative interview. During this interview subjects are asked open-ended questions about their impression of the system and the character.
44	Subjects are emailed with reminders to log into the system to fill out their follow-up questionnaires. When subjects login they fill out one-week recall of activity forms (behavioral follow-up measure) and exercise self-efficacy, decisional balance and stage of change questionnaires.
45+	Subjects who have completed all requirements are debriefed and compensated.

Table 9-1. Experimental Protocol and Measures

9.2.1 Subjects

A preliminary power analysis based on previous studies of therapist experience level and working alliance (Mallinckrodt & Nelson, 1991), working alliance and outcome in cognitive therapy (Safran & Wallner, 1991), and cognitive-behavioral interventions and exercise adoption behavior (Dishman, 1994) indicated that at least 30 subjects per condition were required to achieve statistical significance (based on a one-tailed power analysis, with $\alpha = 0.05$, $\beta = 0.2$). Allowing for 10% attrition in the longitudinal study, this indicated that a total of 99 subjects would be needed.

Subjects were recruited via three advertisements in the MIT school newspaper ("The Tech") and fliers posted around campus (see Appendix B), and directed to a web page that explained the experiment in more detail, provided on-line screening and intake scheduling forms, and randomly assigned subjects to one of the three experimental conditions.

Subjects were screened according to the following criteria:

1. Subjects had to pass the Physical Activity Readiness Questionnaire (PAR-Q) (Chisolm, Collis, Kulak, Davenport, & Gruber, 1975). They had to answer "no" to each of the following questions:
 - Has your doctor ever said that you have a heart condition and that you should only do physical activity recommended by a doctor?
 - Do you feel pain in your chest when you do physical activity?
 - In the past month, have you had chest pain when you were not doing physical activity?
 - Do you lose your balance because of dizziness or do you ever lose consciousness?
 - Do you have a bone or joint problem that could be made worse by a change in your physical activity?
 - Is your doctor currently prescribing drugs (for example, water pills) for you blood pressure or heart condition?
 - Do you know of any other reason why you should not do physical activity?
2. Subjects had to be in contemplation, preparation, or action stages of change with respect to the current CDC/ACSM guidelines for physical activity (Marcus, Selby, Niaura, & Rossi, 1992; Pate et al., 1995). They had to answer "b", "c", or "d" to the following question:
 - Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week. Question: Are you physically active according to that definition?
 - a) Yes, I have been for MORE than 6 months;
 - b) Yes, I have been for LESS than 6 months;
 - c) No, but I intend to in the next 30 days;
 - d) No, but I intend to in the next 6 months;
 - e) No, and I do NOT intend to in the next 6 months.
3. Subjects had to have regular access to a personal computer running Windows that they could install the client software on. They had to answer "yes" to the following question:
 - Do you own a personal computer with the following minimum resources?
 - Microsoft Windows XP, 2000, NT, or 98
 - Pentium II, 233 MHz; 64 MB RAM or better

- Sound card and speakers or headphone
 - 100 MB free hard disk space
4. Subjects had to indicate that they would be available and willing to perform all of the tasks required in the study. They had to respond “very likely” to the following question:
- The compensation for this experiment is a Digimax Digital Pedometer (\$25 value) plus \$25 cash. This experiment requires that you do the following:
 - Come into the Media Lab for a 30 minute interview and orientation session.
 - Install software on your computer (10 minutes).
 - Wear the pedometer at least 5 days per week for a month.
 - Run the software at least 5 days per week for a month (10-15 minutes per session)
 - Two months after the start of the study you will be asked to run the software one final time.
 - You come into the Media Lab for a final 10 minute debrief interview and to get paid.
- How likely would you be to complete these tasks during the Fall semester?

Figure 9-2 shows the flow of subjects through the study. Approximately 986 people visited the recruiting website over a three-week period of time, 13% were screened out

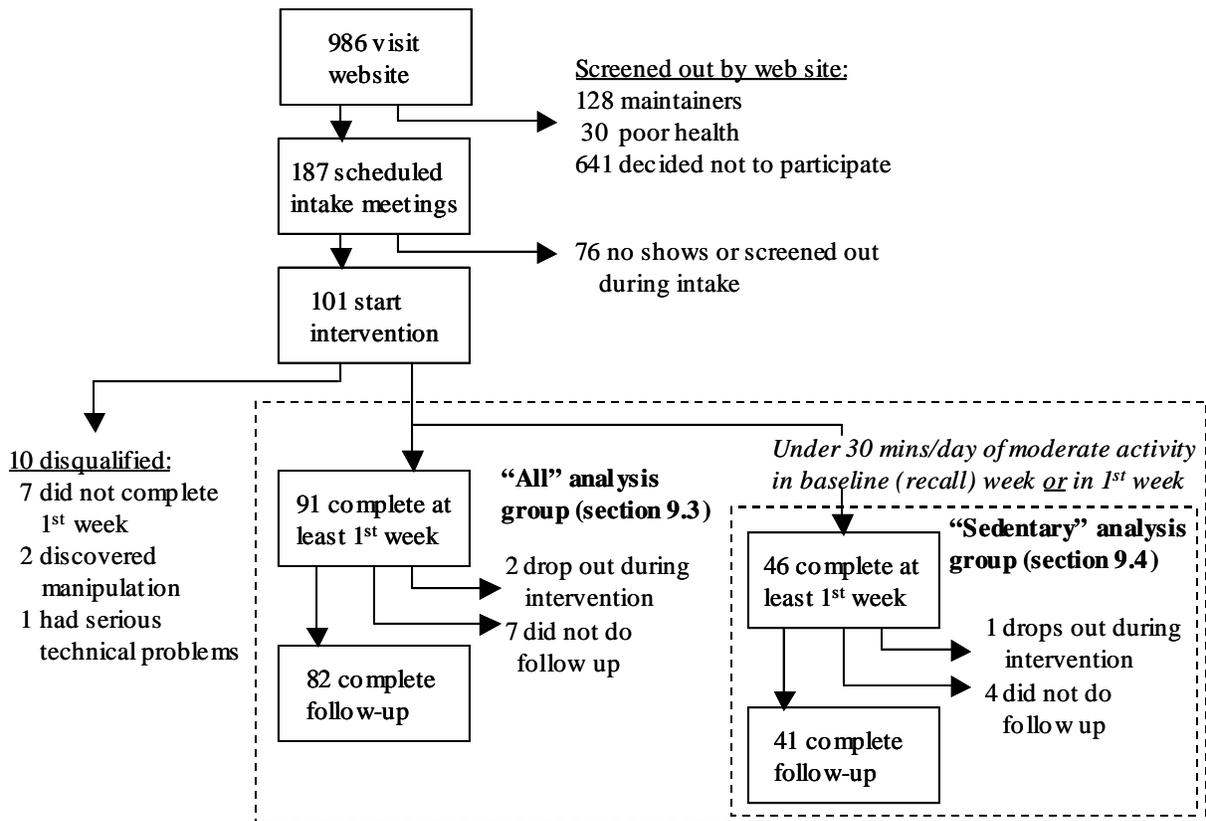


Figure 9-2. Flow of Subjects through Study

because they were already "maintainers" (with respect to exercise stage of change), 3% were screened out due to health reasons, 65% decided to not participate, 19% scheduled intake meetings, and half of those actually showed up, resulting in 101 subjects who completed the intake process. Of these, 10 were later disqualified from data analysis due to non-participation, significant technical problems with the system, or because they made it known to the experimenters that they were aware of the experimental manipulation.

Early into the study it was discovered that 54 of the 91 subjects had reported baseline activity levels over 30 minutes per day, even after correcting for the anomaly described in section 9.2.5. Given that all of the interventions were designed to get sedentary subjects up to this goal level (via educational content, goal setting and dialogue tailored to this end), this group of subjects represented individuals who clearly should not benefit from the intervention as designed. Consequently, a subgroup of the most sedentary subjects was identified for additional analysis. To allow for the fact that some individuals might have had unusually strenuous baseline weeks, this group was defined as those subjects who averaged less than 30 minutes of moderate-or-better activity in their baseline week or in their first week of intervention. Thus, results in the following sections are reported for both "all" 91 subjects (section 9.3) and for the most "sedentary" subjects (section 9.4).

There are several possible reasons why so many subjects had baselines above 30 minutes per day of moderate-or-better activity. First, 25 subjects were in the "action" stage of change, meaning that they were currently getting at least 30 minutes a day of moderate activity, but had been doing this for less than six months. Second, 6 subjects who were actually maintainers were mistakenly admitted to the study (these are included in the "all" analysis group, but excluded from the "sedentary" group). Finally, even though the screening form defined regular physical activity as "any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little," many subjects still interpreted this to mean something closer to vigorous activity (such as running) and thus significantly under-estimated their moderate activity levels on the intake screening instrument (based on discussions with subjects).

For remuneration, subjects were given a digital pedometer to keep (a model MLS 2000 pedometer, manufactured by Yamax and marketed through Walk4Life)--a \$24 value--and \$25 cash upon completion of all tasks in the study.

Breakdown of all subjects by condition and gender is shown in Table 9-1, demographic breakdown of subjects is shown in Table 9-2, and subjects' initial self-assessments of health and fitness is shown in Table 9-3.

		Gender		Total
		female	male	
Condition	CONTROL	16	15	31
	NON-REL	15	12	27
	RELATIONAL	21	12	33
	Total	52	39	91

Table 9-1. Breakdown of All Subjects by Condition and Gender

		Mean	SD	Mean		
Age	Years	24.85	7.40	Education	High school	41%
Height	Inches	66.70	3.99		Bachelors	34%
Weight	Pounds	154.50	36.23		Masters	19%
Body mass index		24.36	4.31		PhD	6%
Ethnicity	Caucasian	54%		Occupation	Student	69%
	Asian	34%			Univ staff	14%
	Hispanic	4%			Industry	6%
	Other	8%			Faculty	4%
Marital	Single	69%			Government	1%
Status	Married	25%			Unemployed	1%
	Domestic Partner	5%			Other	5%

Table 9-2. Demographics for All Subjects

Smoke		3%	Self-rating of sleep	good	24%
Currently engage in regular sport		21%		fair	50%
Previously exercised		89%		poor	26%
Self-rating of overall health	good	42%		Concerned	62%
	fair	57%	Reasons for wanting To increase physical activity	Overall Health	92%
	poor	1%		Lose Weight	55%
	Concerned	97%		Fitness "in and of itself"	58%
about health		Muscle Gain		46%	
Self-rating of diet	good	28%		Appearance	73%
	fair	62%		Anxiety	47%
	poor	10%		Study Better	39%
	Concerned?	86%	Physical activity stage of change	Contemplation	32%
good	7%	Preparation		34%	
fair	66%	Action		27%	
poor	27%	Maintenance		7%	
Self-rating of fitness	poor	100%			
	Concerned				

Table 9-3. Self-assessments of Health at Intake by All Subjects

9.2.2 Apparatus

The experiment makes use of the “FitTrack” client-server software described in chapter 8. The client software is installed on subjects’ home computers and run on a daily basis during intervention. Figure 9-3 shows the basic layout of the client interface. Figure 9-4 shows a screen shot representative of the interface during interactions with forms and educational content (for all conditions), while Figure 9-5 shows a screen shot representative of the interface during interactions with the exercise advisor agent (RELATIONAL and NON-RELATIONAL conditions only).

As mentioned above, subjects were also each given a digital pedometer (a model MLS 2000 pedometer, manufactured by Yamax and marketed through Walk4Life).

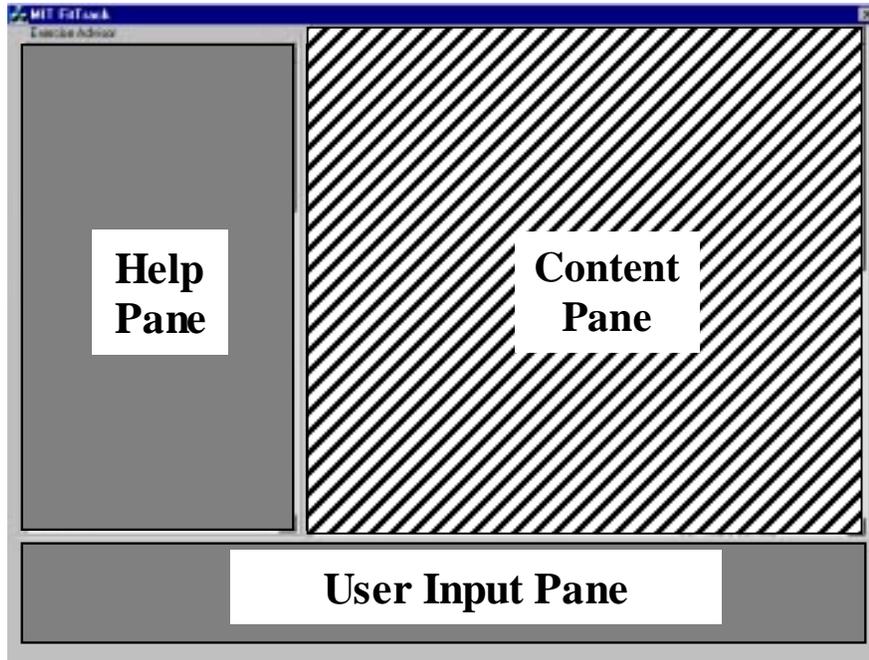


Figure 9-3. Interface Layout of the Client Software

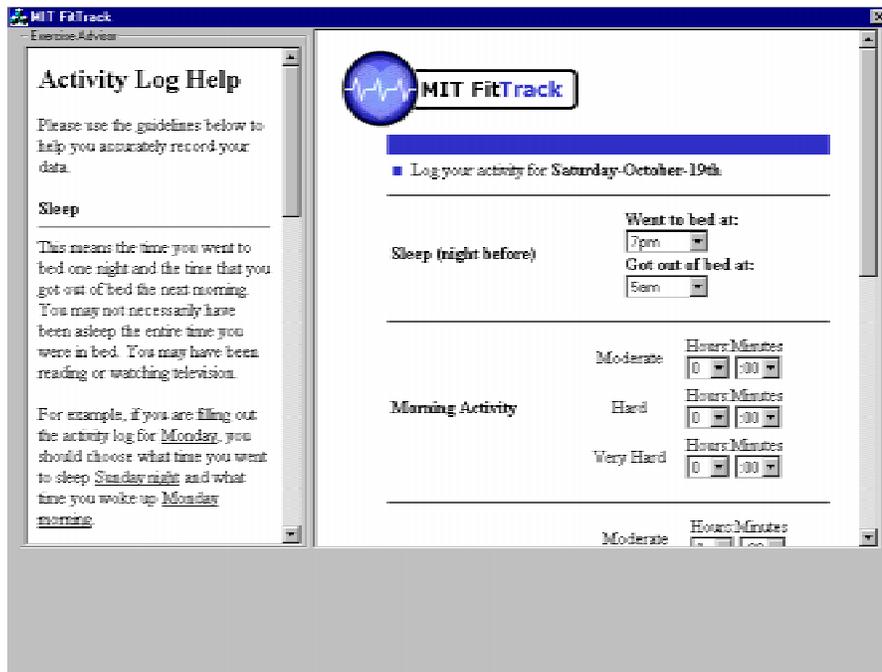


Figure 9-4. Sample Screenshot of Client Software during Form Interaction (all of CONTROL; portions of RELATIONAL/NON-RELATIONAL)

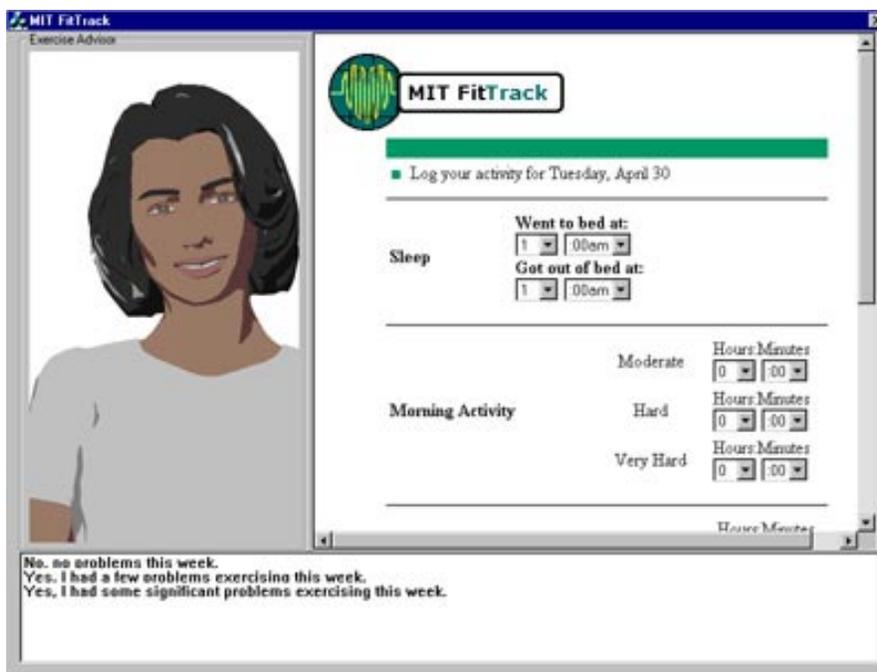


Figure 9-5. Sample Screenshot of Client Software during Agent Interaction

9.2.3 Procedure

During the intake meeting, subjects were first asked to fill out a screening questionnaire. If they did not pass any of the criteria for inclusion in the study (pass the Physical Activity Readiness Questionnaire, in contemplation, preparation or action stage of change, and owned a home computer with the requisite resources for the client software) they were dismissed. Otherwise, they next signed a consent form and filled out demographic and personality questionnaires and asked to provide a password to use in accessing the server. They were then given a pedometer and instructions on using it. If they had a high speed internet connection at home they were given a URL to download the client software from, otherwise they were given a CD-ROM with the client software. Finally, subjects were given an instruction sheet describing how to install and use the software and pedometer, their unique USER-ID (used to identify them in all interactions with the server), and an email address to use if they encountered any problems or had any questions about the study. Subjects were asked to install the software as soon as possible, but wait until they were notified before using their pedometer or logging in for the first time (if they attempted to login before their account was enabled they received an error message instructing them to wait).

Subjects then installed the software on their home computers using a commercial installation program. As part of the process they were prompted for their USER-ID, and this was stored in their system registry so that they did not have to re-enter it every time they logged in (they would still be prompted for their password every session, however, to prevent unauthorized access). Following successful installation, a test program was automatically run for users in NON-RELATIONAL and RELATIONAL conditions to test the functionality of the speech synthesizer and the quality of their sound system. Subjects were told that if the

speech was un-intelligible that they would be provided with headphones (none of the subjects requested this).

Subjects were then notified via email when the server was ready for them to start the study. They were instructed to begin wearing their pedometer and logging in on a daily basis. They were told that it was alright if they could not log in for a day or two, as long as they did not miss more than three consecutive days in a row. Subjects who missed more than three consecutive days were sent daily emails reminding them to log in.

On their first login, baseline physical activity was assessed using a 7-day recall, along with baseline assessments of exercise attitudinal measures (self-efficacy and decisional balance questionnaires).

Every session during the 30-day intervention period had the same structure:

1. Subjects logged into the system (providing their password).
2. They were prompted for estimates of their daily activity and pedometer reading for each day since they last logged in.
3. They were shown self-monitoring charts displaying actual minutes of activity and steps relative to their daily goal level (see Figure 9-6).
4. Subjects were then given any survey questionnaires scheduled for the day of the study (based on days since their first login).
5. For those in RELATIONAL and NON-RELATIONAL conditions only, the left-hand help pane was replaced with the FitTrack logo, Laura walked onto the screen and conducted a 3-5 minute dialogue with them.
6. An educational content page was then displayed. At this point subjects could spend as much time as they liked reading the page, navigating to a table of contents for all educational pages in the library (and viewing any of those pages), viewing their self-monitoring chart again, or using a feedback form to report problems with the software or ask questions.

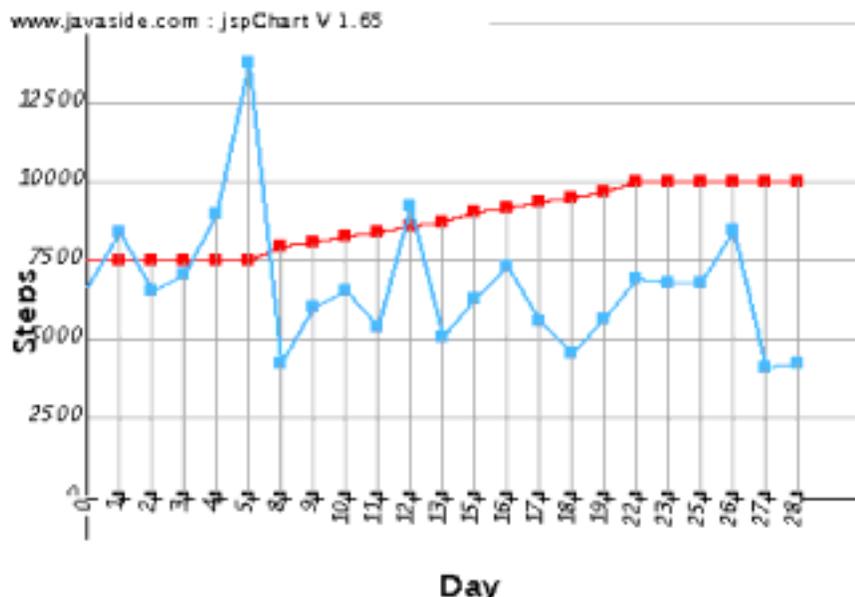


Figure 9-6. Sample Self-Monitoring Chart. Blue (jagged) line represents actual steps walked; red line represents daily goal.

After finishing the last day of intervention, subjects were instructed (via the client software) to wait until they heard from FitTrack before logging in again. They were not given any additional instructions on what to do during the relapse period. A subset of subjects were contacted by to participate in a 5-minute qualitative interview about their experience.

Six weeks after the day they first logged in, subjects were notified via email to log in one final time. In this session they were given a final 7-day recall of activity (to check for extinction of any gains in physical activity), as well as other attitudinal, relational and system evaluation measures. They were then instructed to contact the experimenters to make an appointment for debriefing and payment.

At the final debrief meeting, subjects were told about the goals and design of the study and which condition they were in, and given a voucher for compensation for their participation.

9.2.4 Measures

In addition to demographic questionnaires, system evaluation questionnaires and qualitative interviews, the quantitative measures used in the study fall into four groups, corresponding to the four sets of hypotheses described above.

Relational Measures

- The Working Alliance Inventory (WAI) (Horvath & Greenberg, 1989), a self-report measure used to assess the agent-subject relationship. This measure was slightly modified for exercise adoption and use with an animated character (see Appendix B). This instrument is a 36-item questionnaire comprised of three subscales: bond (WAI/BOND) – the degree to which the helper and helpee like and trust each other; task (WAI/TASK)—the degree to which the helper and helpee agree on the therapeutic tasks to be performed; and goal (WAI/GOAL)—the degree to which the helper and helpee agree on the goals of therapy. The WAI was administered on days 7 and 27 of the intervention.
- Single item questions on relationship with Laura. These include: “How much do you like Laura?” (LIKE LAURA); “How would you characterize your relationship with Laura?” (CLOSENESS, ranging from “Complete Stranger” to “Close Friend”); “How useful were your discussions with Laura?” (USEFUL LAURA) and “How much would you like to continue working with Laura?” (CONTINUE LAURA). These questions were asked on day 30 of the intervention, and CONTINUE LAURA was asked again as part of the follow up.
- Sentimental farewell (FAREWELL). At the end of the last interaction with Laura (day 30), subjects were given the opportunity to choose a brief farewell (“Bye.”) or a “sentimental” farewell (“Take care Laura, I’ll miss you.”). This behavioral measure tracked whether each subject chose the sentimental version or not, under the assumption that a closer bond would lead to an increased frequency of sentimental partings.
- A multiple choice question “Who was most helpful in getting you to exercise over the last month?”, with possible responses being “Laura”, “friends”, “family”, “workout buddy” or “none of the above” asked on day 30 of the intervention.

Physical Activity Measures

- A baseline measure of physical activity was measured by a 7-day recall on day 1 of the intervention, following the form and procedure in (Sallis, 1997), aggregated into average minutes per day of moderate or better activity (MIN/DAY), and days per week at or above goal (DAY/WK>30MIN, 30 minutes of moderate-or-better activity).
- Daily measures of physical activity were measured by a 1-day recall on days 1-30 of the intervention, following the form and procedure in (Sallis, 1997) (except given in a 1-day format), aggregated into average minutes per day of moderate or better activity (MIN/DAY), days per week at or above goal (DAY/WK>30MIN), and slopes of minutes per day change (MIN/DAY CHNG) and days per week over goal change (D/WK>30M CHNG) per week of the study.
- Daily step count was measured by digital pedometer and reported by subjects via a web form on days 1-30 of the intervention. As mentioned above, a model MLS 2000 pedometer, manufactured by Yamax and marketed through Walk4Life was used. Yamax pedometers have been used in a number of prior studies of physical activity and have been demonstrated to have high reliability (Bassett et al., 1996; Bassett, Cureton, & Ainsworth, 2000). Raw step counts were converted into average steps per day aggregated per week (STEP/DAY), days per week over goal (10,000 steps; DAY/WK>10KSTEP), and slopes of steps per day change (STEP/DAY CHNG) and days per week over goal change (D/WK>10KS CHNG) per week of the study.
- Physical activity at follow up was measured by a 7-day recall on day 44 of the intervention, following the form and procedure in (Sallis, 1997), aggregated into average minutes per day of moderate or better activity (MIN/DAY), and days per week at or above goal (30 minutes; DAY/WK>30MIN).

Participation Measures

- Attrition. The number of subjects in each condition who excused themselves from the study or stopped logging in and responding to emails was tracked.
- Days logged in (DAYS). Since subjects were not required to log in every day, the actual number of days they logged in was tracked as a behavioral measure of participation.
- Total number of logins (LOGINS). Subjects could login multiple times per day if desired, in order to read through the educational library contents and view their self-monitoring charts (although the agent would only appear during the first login each day for subjects in RELATIONAL and NON-RELATIONAL conditions).
- Pages per session (PAGES). Subjects were free to read through as many pages of educational content as they liked. The number of pages viewed per day was tracked as a behavioral participation measure.

Exercise Attitudinal Measures

- Exercise Stage-of-Change (STAGE), administered at intake, the last day of the intervention, and at follow up (Marcus et al., 1992). Stage of change is based on on a single question about exercise behavior.

- Exercise self-efficacy (SELF-EFFICACY), administered at intake, the last day of the intervention, and at follow up (Marcus et al., 1992). Self-efficacy is based on an assessment of ability to exercise given 18 different obstacles.
- Exercise decisional-balance, administered at intake, the last day of the intervention, and at follow up (Nigg, Rossi, Norman, & Benisovich, 1998). Decisional balance is based on the number of advantages (PROS) vs. disadvantages (CONS) to exercise that a subject agrees with (the questionnaire has 5 of each).

9.2.5 Anomalies

Three days after subjects started interacting with the system it was discovered that many of them had entered unusually high values for their activity time estimates. After determining that the problem was partly due to misunderstandings regarding the definition of "Moderate Activity", a new definition was provided on the form (consistent with (Sallis, 1997)) and the 34 subjects who had already entered activity estimates were asked to re-do their estimates.

However, even with the new definition, 51 subjects were found to have baselines greater than 30 minutes per day (based on their 7-day recall data). Consequently, a feature was added to the system so that users who had been performing consistently at or above 30 minutes a day for two weeks were given the option of raising their goal level of exercise to 45, 60, 75 or 90 minutes a day of moderate or better activity. This was implemented in such a way that the feature appeared for all subjects who met the criteria on the same day of the study (day 8). Only 26 of the 46 eligible subjects took advantage of this feature (distributed roughly equally among the study conditions), and only half of these (13) set new goals that were actually above their baseline levels of activity.

9.3 Quantitative Results for All Subjects

Quantitative results were analyzed as follows. The hypotheses listed above are evaluated using one-tailed, planned comparisons between RELATIONAL and NON-RELATIONAL groups and between groups with the agent (RELATIONAL and NON-RELATIONAL together) and without it (CONTROL). Measurements taken at more than one time point are analyzed using dependent-means t-tests on all subjects and each treatment group individually, to determine if significant changes occurred. Error bars on plots of means correspond to 90% confidence intervals.

Table 9-4 summarize the results of the between-group planned comparisons, while Table 9-5 summarizes the results of the dependent-means t-tests for changes over time.

9.3.1 Relational Measures for All Subjects

Working Alliance Inventory

Figure 9-7 shows the means for the two administrations of the working alliance inventory questionnaire. The only significant differences are on the bond subscales of both surveys, in the hypothesized direction (greater for RELATIONAL): for day 7, $t(58)=1.75$, $p<.05$; and for day 27, $t(57)=2.26$, $p<.05$. The fact that there was no significant differences in the task and goal subscales (and the composite measure) is not too surprising, given that little effort was made in the design of the RELATIONAL strategies to address clarification of and agreement on tasks and goals.

Measure	Day	CONTROL		NON-REL		RELATIONL		REL>NON-REL			AGENT>CONTRL		
		Mean	SD	Mean	SD	Mean	SD	df	t	p	df	t	p
WAI/COMP	7			4.798	0.821	4.863	0.659	58	1.094	0.139			
	27			4.767	0.912	4.902	0.800	57	1.194	0.119			
WAI/BOND	7			4.297	0.926	4.506	0.804	58	1.746	0.043			
	27			4.333	0.946	4.638	1.000	57	2.263	0.014			
WAI/TASK	7			5.127	0.929	5.273	0.653	58	1.322	0.096			
	27			5.109	0.996	5.213	0.860	57	0.586	0.280			
WAI/GOAL	7			4.971	0.839	4.810	0.888	58	0.239	0.406			
	27			4.859	0.985	4.856	0.933	57	0.212	0.417			
LIKE LAURA	30			4.609	1.305	5.207	1.346	57	2.035	0.023			
CLOSENESS	30			2.261	0.752	2.517	0.829	57	1.621	0.055			
USEFUL LAURA	30			2.348	0.982	2.621	0.979	57	1.256	0.107			
CONTINUE LAURA	30			2.043	0.928	2.517	0.949	57	2.429	0.009			
	44			2.043	0.878	2.621	1.049	53	1.828	0.037			
FAREWELL	30			0.348	0.487	0.690	0.471	54	2.796	0.004			
MIN/DAY	-6-0	50.51	41.92	40.24	33.44	54.92	75.51	58	0.897	0.187	88	0.054	0.478
	1-7	41.37	20.3	41.9	19.07	40.11	17.79	58	0.076	0.470	88	0.875	0.192
	8-14	37.54	19.1	39.94	23.45	37.2	17.12	58	0.099	0.461	88	0.129	0.449
	15-21	40.57	19.66	42.62	20.79	39.26	15.28	58	0.822	0.207	87	0.380	0.353
	22-30	39.08	22.21	41.09	19.2	38.86	18.2	57	0.166	0.434	86	0.027	0.489
	38-44	27.49	12.58	34.26	19.81	32.35	26.34	53	0.201	0.421	81	0.380	0.353
DAY/WK>30MIN	-6-0	4.08	2.798	3.542	2.431	3.741	2.782	58	0.871	0.194	88	0.061	0.476
	1-7	4.32	2.096	4.417	1.586	4.185	1.733	58	0.236	0.407	88	0.266	0.395
	8-14	4.64	2.325	4.375	1.996	4.481	2.19	58	0.683	0.249	88	0.005	0.498
	15-21	4.36	2.177	5.125	2.007	4.593	1.886	58	0.930	0.178	87	1.058	0.146
	22-30	5.32	2.854	6.25	2.541	6.222	2.407	57	0.244	0.404	86	1.536	0.064
	38-44	3.08	1.998	3.875	2.29	3.667	2.449	53	0.010	0.496	79	1.344	0.091
STEP/DAY	1-7	8242	2654	9425	2891	8800	3359	58	0.336	0.369	88	1.145	0.128
	8-14	8869	2998	9926	3343	9414	3796	58	0.573	0.284	88	0.764	0.223
	15-21	9709	3291	10208	3025	10091	3031	57	0.450	0.327	86	0.105	0.458
	22-30	9052	3890	10435	3597	9523	3277	57	0.983	0.165	86	0.600	0.275
DAY/WK>10KSTEP	1-7	2.04	1.791	2.25	1.539	2.519	1.949	55	0.777	0.220	84	0.661	0.255
	8-14	2.12	1.922	3.208	2.303	2.667	2.304	55	0.746	0.229	84	1.050	0.148
	15-21	2.76	2.006	3.458	2.043	3.259	2.105	55	0.161	0.436	83	0.421	0.337
	22-30	2.68	2.626	3.958	2.805	3.556	2.455	56	0.654	0.258	84	1.542	0.063
MIN/DAY CHNG	1-30	-0.385	4.449	0.025	5.245	-0.168	6.538	57	0.192	0.424	86	0.854	0.198
D/WK>30M CHNG	1-30	0.272	0.693	0.625	0.786	0.622	0.796	57	0.229	0.410	86	1.982	0.025
STEP/DAY CHNG	1-30	326.9	1219	331.2	750.4	284.8	1108	57	0.652	0.258	86	0.207	0.418
D/WK>10KS CHNG	1-30	0.256	0.764	0.538	0.741	0.37	0.675	53	0.860	0.197	79	1.295	0.100
DAYS	1-30	19.32	5.288	19.15	5.934	17.58	5.64	58	1.049	0.149	88	0.772	0.221
LOGINS	1-30	24.71	9.253	24.11	7.885	22.45	7.005	58	0.861	0.196	88	0.796	0.214
PAGES	1-30	1.066	0.077	1.159	0.232	1.389	0.891	58	1.306	0.098	88	1.701	0.046
STAGE	Intake	2.074	0.958	2.080	0.909	2.067	0.944	58	0.224	0.412	88	0.314	0.377
	30	2.852	0.989	2.880	0.726	2.867	0.730	57	0.076	0.470	86	0.101	0.460
	44	2.630	1.079	2.800	0.816	2.700	0.837	53	0.446	0.329	79	0.557	0.290
SELF-EFFICACY	1	2.453	0.757	2.320	0.658	2.307	0.586	58	0.634	0.264	88	0.700	0.243
	29	2.556	0.714	2.633	0.593	2.733	0.555	57	0.402	0.345	86	0.961	0.170
	44	2.671	0.794	2.916	0.572	2.741	0.567	53	1.134	0.131	79	1.026	0.154
PROS	1	2.948	0.669	3.136	0.709	3.227	0.614	58	0.681	0.249	88	1.540	0.064
	29	2.830	0.688	3.088	0.648	3.053	0.714	57	0.128	0.449	86	1.396	0.083
	44	2.911	0.757	3.112	0.638	3.020	0.690	53	0.510	0.306	79	0.943	0.174
CONS	1	3.385	0.592	3.192	0.649	3.387	0.664	58	0.941	0.175	88	0.570	0.285
	29	3.289	0.564	3.240	0.676	3.280	0.616	57	0.135	0.447	86	0.321	0.375
	44	3.319	0.601	3.288	0.643	3.313	0.577	53	0.154	0.439	79	0.125	0.450
CONTINUE FT	30	2.926	0.675	2.920	0.812	3.000	0.830	57	0.976	0.167	86	0.130	0.449
	44	2.704	0.869	2.760	0.879	3.000	0.788	53	1.067	0.145	79	0.888	0.189

Table 9-4. Between-group Planned Comparisons for All Subjects

Measure	Change		ALL CONDS			CONTROL			NON-REL			RELATIONL		
	From	To	df	t	p	df	t	p	df	t	p	df	t	p
	Day1	Day2												
WAI/COMP	7	27	54	0.205	0.838				24	0.014	0.989	29	0.361	0.720
WAI/BOND	7	27	54	0.519	0.606				24	0.376	0.710	29	1.489	0.147
WAI/TASK	7	27	54	0.134	0.894				24	0.409	0.686	29	0.661	0.514
WAI/GOAL	7	27	54	0.155	0.877				24	0.081	0.936	29	0.329	0.745
CONTINUE LAURA	30	44	54	0.868	0.389				24	0.625	0.538	29	0.619	0.541
MIN/DAY	-6-0	22-30	81	1.470	0.145	26	1.274	0.214	24	0.124	0.903	29	1.104	0.279
	1-7	22-30	81	0.691	0.492	26	0.758	0.456	24	0.109	0.914	29	0.358	0.723
	22-30	38-44	81	3.626	0.001	26	2.480	0.020	24	1.959	0.062	29	1.804	0.082
DAY/WK>30MIN	-6-0	22-30	81	6.653	0.000	26	2.323	0.028	24	5.284	0.000	29	4.347	0.000
	1-7	22-30	81	6.272	0.000	26	2.401	0.024	24	3.818	0.001	29	4.597	0.000
	22-30	38-44	81	8.990	0.000	26	4.043	0.000	24	5.322	0.000	29	6.530	0.000
STEP/DAY	1-7	22-30	81	1.778	0.079	26	1.197	0.242	24	2.366	0.026	29	0.236	0.815
DAY/WK>10KSTEP	1-7	22-30	77	3.986	0.000	25	1.355	0.188	23	3.591	0.002	27	2.055	0.050
STAGE	Intake	30	81	6.988	0.000	26	3.403	0.002	24	4.000	0.001	29	4.738	0.000
	30	44	81	2.019	0.047	26	1.185	0.247	24	1.000	0.327	29	1.409	0.169
SELF-EFFICACY	1	29	81	4.782	0.000	26	0.872	0.391	24	3.314	0.003	29	4.750	0.000
	29	44	81	2.770	0.007	26	1.525	0.139	24	4.550	0.000	29	0.085	0.933
PROS	1	29	81	1.998	0.049	26	1.418	0.168	24	0.456	0.653	29	1.540	0.134
	29	44	81	0.393	0.695	26	1.147	0.262	24	0.225	0.824	29	0.308	0.760
CONS	1	29	81	0.902	0.370	26	1.124	0.271	24	0.499	0.622	29	0.823	0.417
	29	44	81	0.740	0.462	26	0.386	0.703	24	0.611	0.547	29	0.339	0.737
CONTINUE FT	30	44	81	1.520	0.133	26	1.442	0.161	24	1.163	0.256	29	0.000	1.000

Table 9-5. Longitudinal Changes for All Subjects

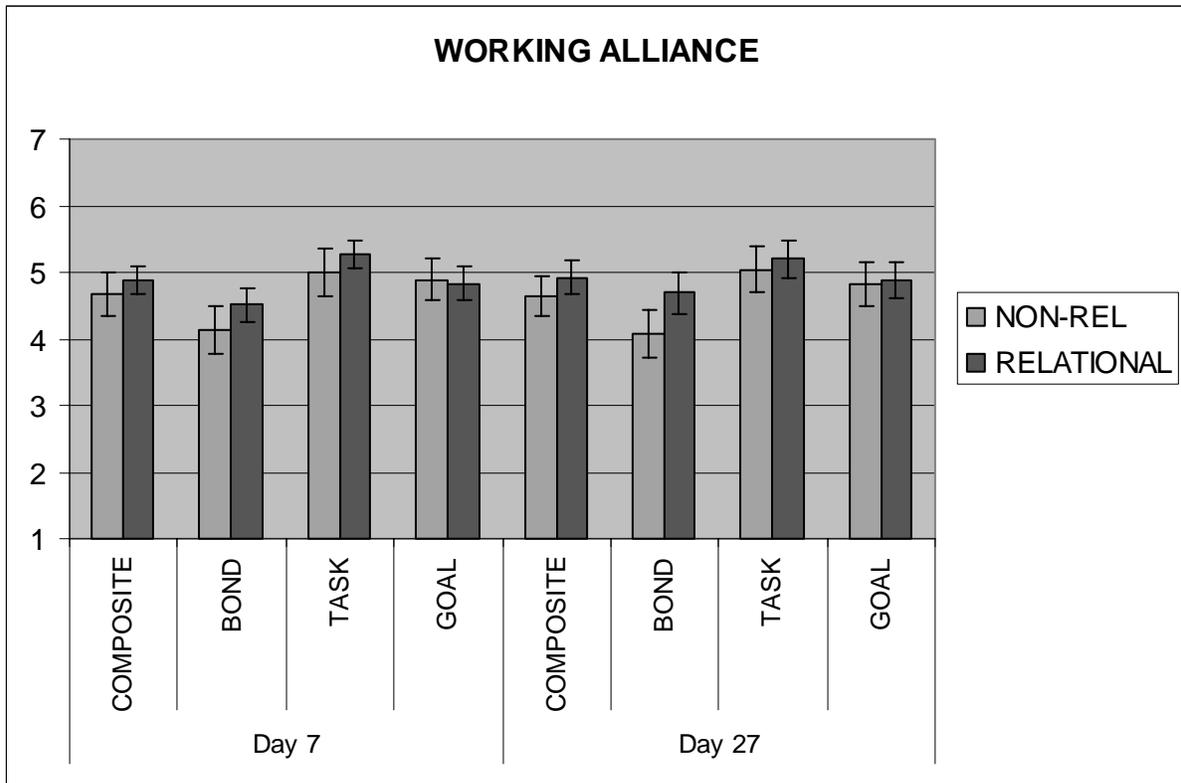


Figure 9-7. Working Alliance Inventory Results for All Subjects

Single Item Relational Questions

In response to the question “How much do you like Laura?”, subjects in RELATIONAL condition reported that they liked her significantly more than those in the NON-RELATIONAL group, $t(57)=2.04$, $p<.05$, see Figure 9-8.

Subjects in RELATIONAL condition also reported a closer relationship with Laura (“How would you characterize your relationship with Laura?”), $t(57)=1.62$, $p=.06$, approaching significance.

There was no significant difference between RELATIONAL and NON-RELATIONAL in how useful they thought their discussions with Laura were.

When asked at the end of the intervention period and again at follow up if they would like to continue working with Laura, subjects in the RELATIONAL condition responded much more favorably compared with the NON-RELATIONAL group, $t(57)=2.43$, $p=.009$ and $t(53)=1.83$, $p<.05$, respectively (see Figure 9-9).

Sentimental Farewell

Given the opportunity to give Laura a sentimental farewell at the end of the intervention period, significantly more subjects in the RELATIONAL group took this option (69%) than in the NON-RELATIONAL condition (35%), $t(54)=2.80$, $p=.004$, see Figure 9-10.

Helpfulness of Laura

Figure 9-11 shows the results of asking subjects about who had been the most helpful in getting them to exercise over the intervention period. The “None of the Above” category is problematic, since it represents the cases in which the subject thought they helped themselves

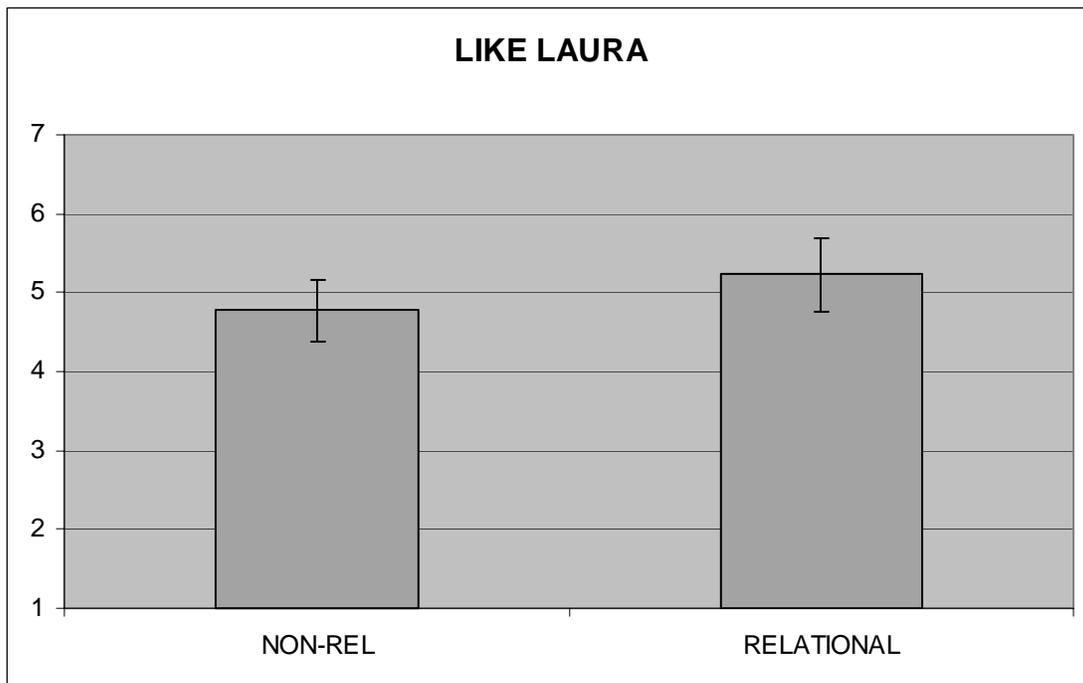


Figure 9-8. Reported LIKING of Laura by All Subjects

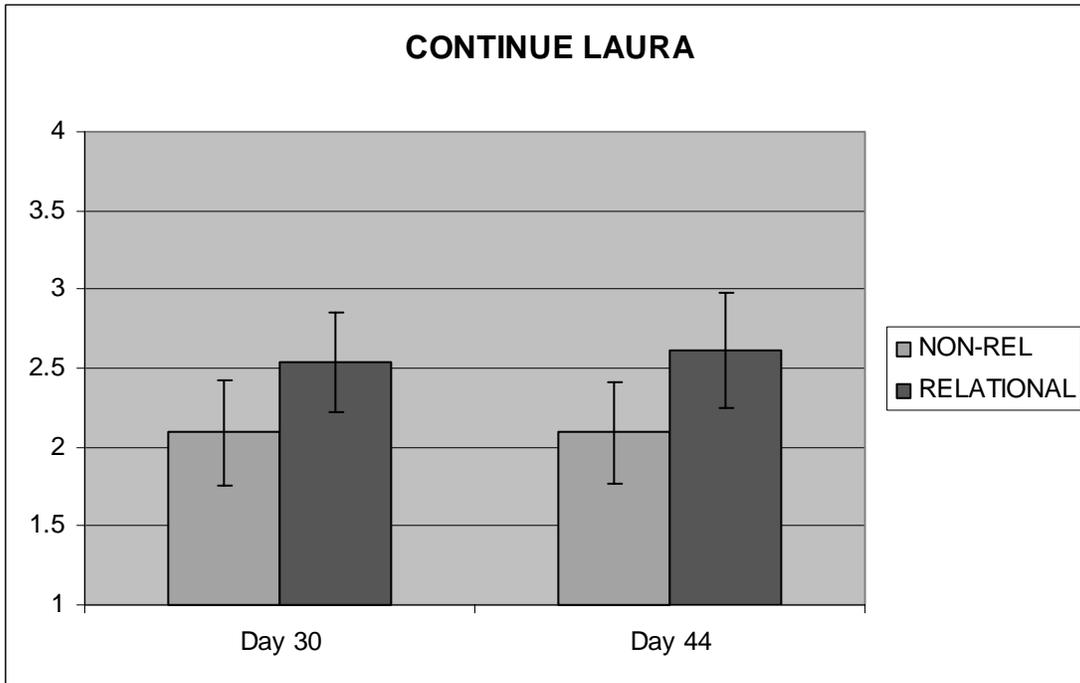


Figure 9-9. Reported Desire to Continuing Working with Laura by All Subjects

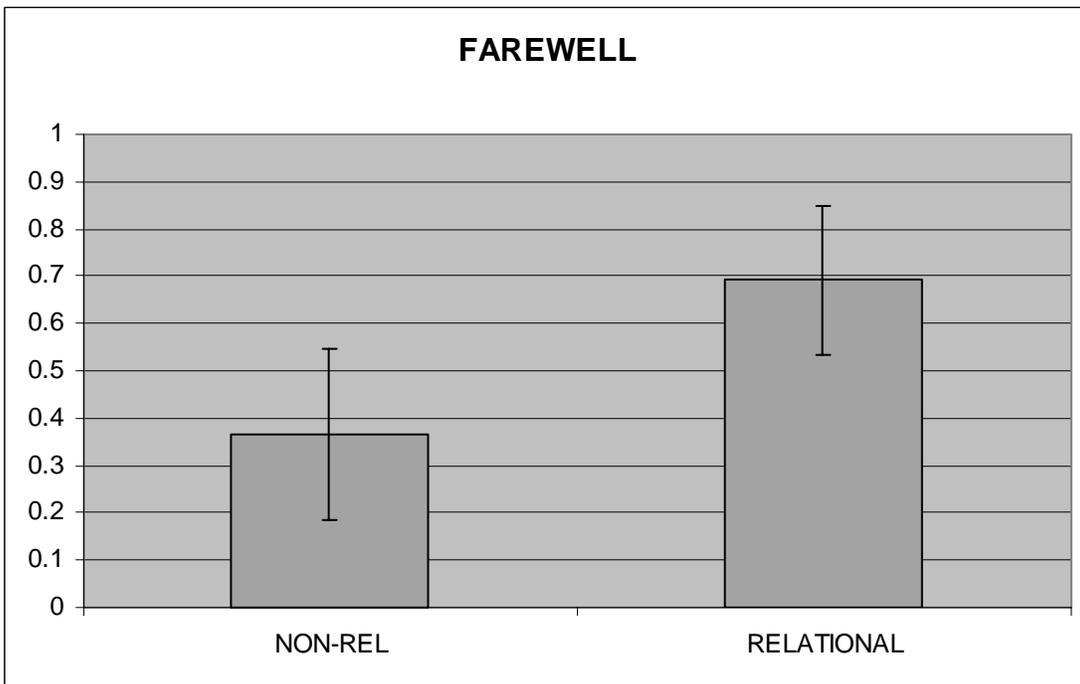


Figure 9-10. Frequency of Using a Sentimental Farewell in Last Interaction
 most, another person not listed helped them most, or if they felt that no-one helped them most.

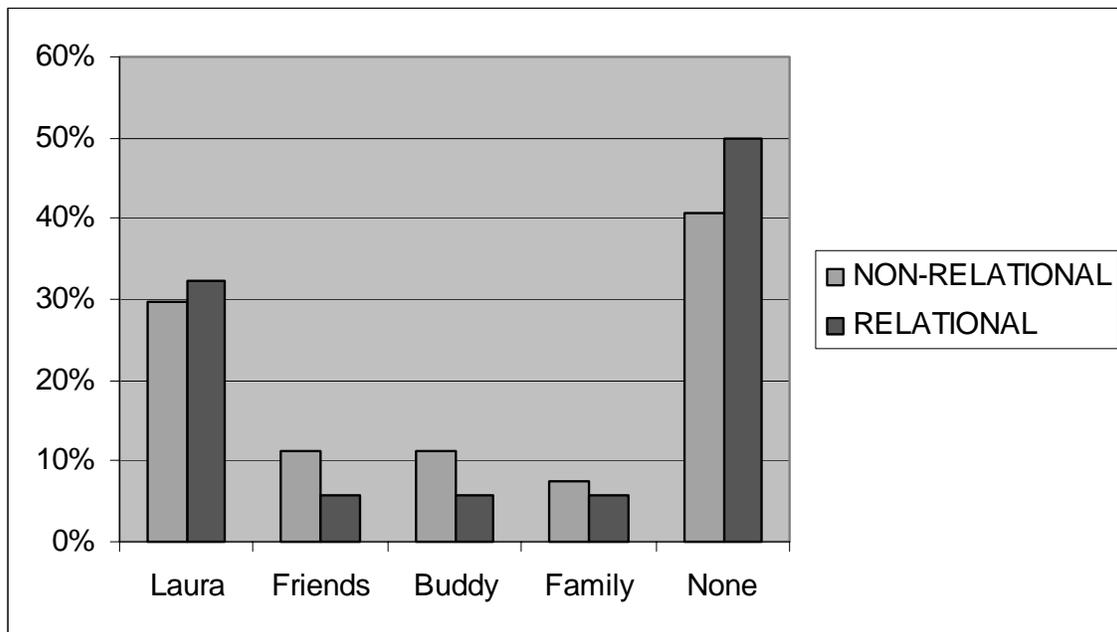


Figure 9-11. Who was Most Helpful in Getting Subjects to Exercise

Thus, excluding this category from analysis, significantly more subjects said that Laura helped them than friends, family or their workout buddy, $X^2(df=3, n=41) = 11.19, p<.05$.

Evaluation of Hypothesis H1

Overall, hypothesis H1 received strong support; subjects in the RELATIONAL condition felt that they had more reciprocal trust, respect and appreciation with Laura (on the bond dimension of the working alliance inventory and the single item question about liking), that they liked Laura more, and that they would like to continue working with her relative to subjects in the NON-RELATIONAL group. And, at the end of their relationship with her, significantly more subjects in the RELATIONAL condition chose a sentimental farewell than those in the NON-RELATIONAL group.

9.3.2 Physical Activity Measures for All Subjects

Minutes per Day of Moderate-or-Better Activity

There were no significant differences between the planned comparison groups for any week for self-reported minutes of moderate-or-better physical activity per day. The only significant change over time was a drop between the last week of the intervention and follow up for all groups combined, paired $t(81)=3.63, p=.001$, as well as for the CONTROL group alone, paired $t(26)=2.48 p<.02$.

Days per Week Meeting Time Goal (30 minutes/day moderate-or-better activity)

There were no significant differences between the planned comparison groups for any week for the DAY/WK>30MIN measure. All groups combined improved significantly during the intervention period, paired $t(81)=6.27, p<.001$, and decreased significantly between the

end of intervention and follow up, paired $t(81)=8.99$ $p<.001$, as did each of the groups independently (see Table 9-5 for details).

Steps per Day

There were no significant differences between the planned comparison groups for any week for number of steps reported by the pedometers. The only significant longitudinal change in this measure was a significant increase during the intervention period for subjects in the NON-RELATIONAL condition, paired $t(24)=2.34$ $p<.05$.

Days per Week Meeting Step Goal (10,000 steps/day)

There were no significant differences between the planned comparison groups for any week for the DAY/WK>10KSTEP measure. All groups combined improved significantly during the intervention period, paired $t(77)=3.99$ $p<.001$, as did subjects in the RELATIONAL and NON-RELATIONAL groups (see Table 9-5 for details).

Rate of Change

Figure 9-12 shows the slopes of least-squares lines fit to each of the four physical activity measures over the four weeks of the intervention (note that slope for steps has been normalized to 1,000 steps per day per week of the study for comparison purposes). Planned comparison analyses indicate that there is a significant difference between the CONTROL group and the other two groups in the rate of change in the DAY/WK>30MIN measure, $t(86)=1.98$ $p<.05$.

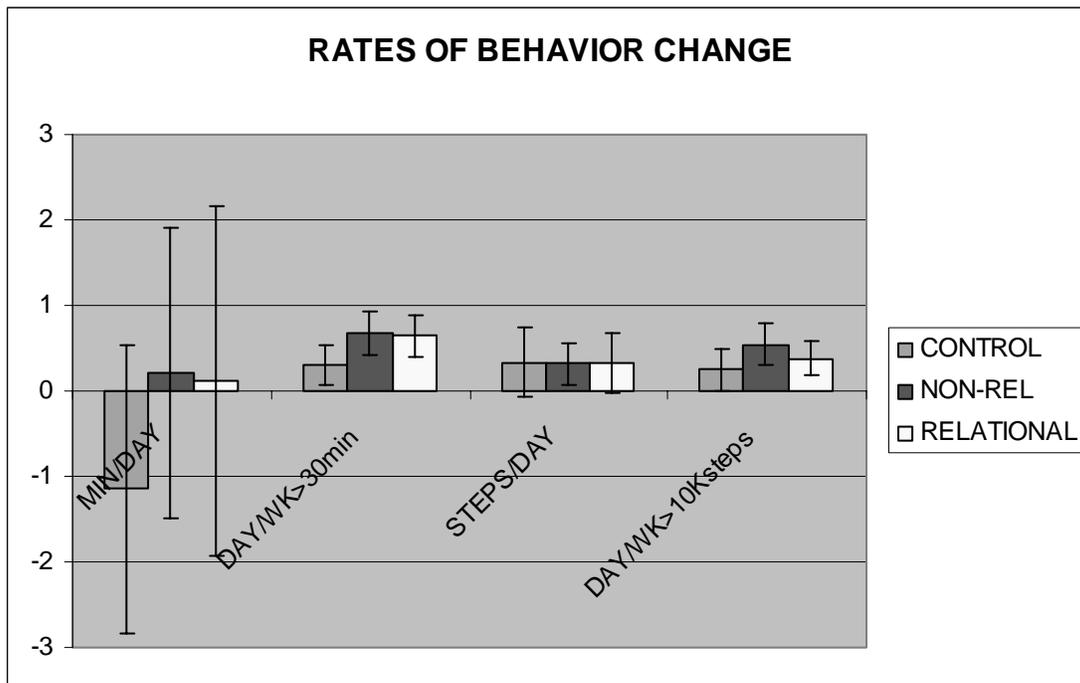


Figure 9-12. Rate of Behavior Change per Week of Intervention

Evaluation of Hypothesis H2

There was no support of hypothesis H2 (a), that subjects in RELATIONAL condition would outperform those in NON-RELATIONAL condition. However, the data suggest that H2(b) may be supported; the groups with the agent outperformed the CONTROL group when looking at the rate of change in days per week over 30 minutes of moderate-or-better activity.

9.3.3 Participation Measures for All Subjects

Attrition

Fortunately for most of the study measures, 82 of the 91 subjects who started the study (logged in at least 5 times) continued through to follow up. Unfortunately, however, this provided little information about distribution of attrition rates as a measure of participation. A software problem was responsible for two of the subjects not completing the follow up surveys (reminder emails did not get sent out to them). One of the subjects who dropped out during intervention contacted the experimenter and said she unexpectedly had to go out of the country for a length of time. Otherwise, reasons for dropping were not given; the other 6 subjects just stopped logging in and responding to email. Table 9-6 shows the breakdown by condition of the 7 subjects who voluntarily dropped from the study.

When	CONTROL	NON-REL	RELATIONAL
Intervention	1	0	1
Follow Up	2	3	0

Table 9-6. Attrition Rates by Condition

Logins

There were no significant differences between treatment groups with respect to either number of days logged in, or total number of logins.

Pages Viewed per Session

Figure 9-13 shows the number of educational content pages viewed by subjects. The results of the planned comparisons indicate that there is significant difference in pages viewed by those in the CONTROL group compared with the other two groups, $t(88)=1.70$, $p<.05$.

Evaluation of Hypothesis H3

There was no support for hypothesis H3(a), that subjects in the RELATIONAL group would show more signs of participation in the FitTrack program than those in NON-RELATIONAL. However, the data suggests some support of hypothesis H3(b), in that subjects in RELATIONAL and NON-RELATIONAL together viewed more education content pages than those in the CONTROL group.

9.3.4 Exercise Attitudinal Measures for All Subjects

Exercise Stage of Change

There were no significant differences between groups in stage of change at any time, but there was a significant increase in stage of change between intake and the end of intervention

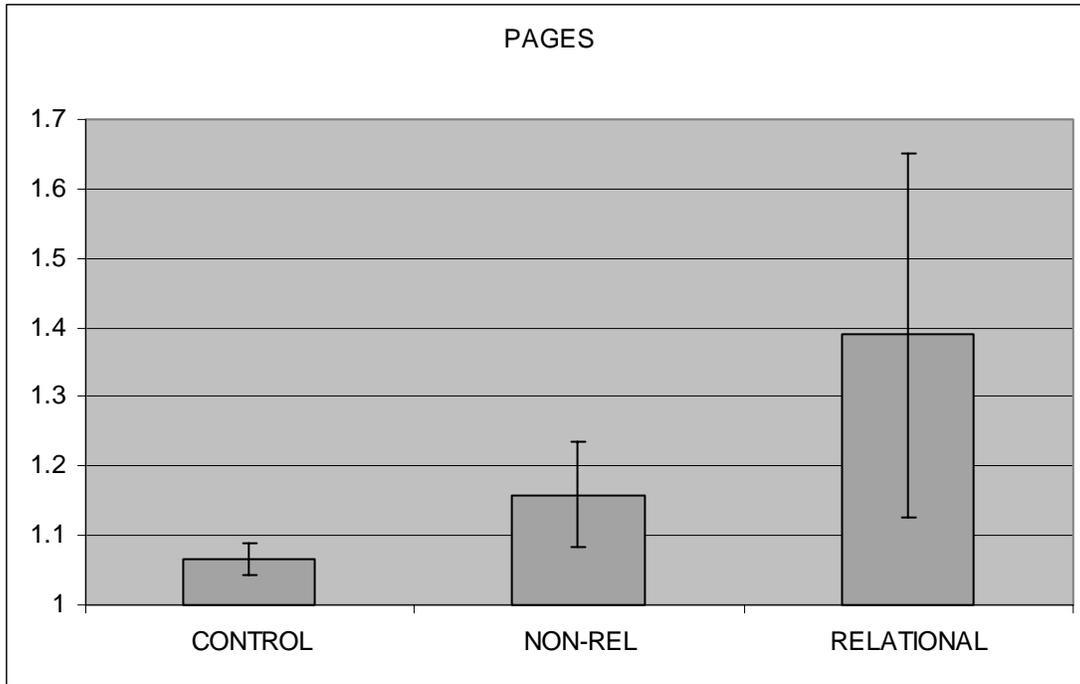


Figure 9-13. Educational Content Pages Viewed per Session

for all subjects combined, paired $t(81)=6.99$ $p<.001$, as well as for each group individually. There was also a significant decrease in stage of change between the end of intervention and follow up for all groups combined, paired $t(81)=2.02$ $p<.05$. Table 9-7 shows the transitions of subjects among stages between intake and the end of intervention.

		At End of Intervention				Total
		Contemplation	Preparation	Action	Maintenance	
At Intake	Contemplation	5	3	20	0	28
	Preparation	2	1	24	4	31
	Action	1	2	15	6	24
	Maintenance	1	0	1	4	6
Total		9	6	60	14	89

Table 9-7. Stage of Change Transitions for All Subjects

Exercise Self-Efficacy

There were no significant differences between groups in self-efficacy at any time. However, there was a significant increase in self-efficacy between intake and the end of intervention, paired $t(81)=4.78$ $p<.001$, and between the end of intervention and follow up, paired $t(81)=2.77$ $p=.007$, for all groups combined.

Exercise Decisional Balance

There were no significant differences between groups in either the advantages of exercise (“pros”) or disadvantages of exercise (“cons”), however, there was a significant decrease in

“pros” for all groups combined between intake and the end of intervention, paired $t(81)=2.00$ $p<.05$.

Evaluation of Hypothesis H4

There was no support for hypotheses H4 (a) or H4 (b); there were no significant differences between groups with respect to attitudes towards exercise, according to the planned comparisons.

9.3.5 Other Quantitative Measures for All Subjects

System Evaluation

At the end of the intervention, subjects were asked to evaluate the FitTrack system and its components. Table 9-8 summarizes the responses to questions about how useful various parts of the system were. There were no significant differences between groups on any of these measures with respect to the planned comparisons. When asked to what degree they would like to continue working with the FitTrack system at the end of intervention, and again at follow up, 24% replied “Very much”, 45% said “Somewhat”, 28% said “Slightly” and only 3% said “Not at all”. There were no significant differences between groups on these evaluations or on changes in these evaluations over time.

	Educational Content	Monitoring Charts	Laura
Extremely	4%	25%	2%
Very	22%	32%	13%
Moderately	37%	30%	34%
Slightly	34%	10%	34%
Not at all	3%	3%	18%

Table 9-8. Usefulness of FitTrack System Features

9.4 Quantitative Results for Sedentary Subjects

As discussed in section 9.2.1, the above analyses were repeated for a subgroup of the most sedentary subjects, defined as those who averaged less than 30 minutes of moderate-or-better activity in their baseline week or in their first week of intervention.

The breakdown of this group by condition is shown in Table 9-5. Table 9-6 summarize the results of the between-group planned comparisons, while Table 9-7 summarizes the results of the dependent-means t-tests for changes over time.

		Gender		Total
		female	male	
Condition	CONTROL	12	4	16
	NON-RELATIONAL	9	2	11
	RELATIONAL	12	7	19
Total		33	13	46

Table 9-5. Sedentary Subjects by Condition and Gender

Measure	Day	CONTROL		NON-REL		RELATIONL		REL>NON-REL			AGENT>CONTRL		
		Mean	SD	Mean	SD	Mean	SD	df	t	p	df	t	p
WAI/COMP	7			4.958	0.713	4.98	0.649	28	0.867	0.197			
	27			4.517	1.002	5.09	0.788	28	2.153	0.020			
WAI/BOND	7			4.208	0.635	4.764	0.71	28	2.551	0.008			
	27			3.938	0.875	4.847	1.024	28	3.458	0.001			
WAI/TASK	7			5.479	0.837	5.31	0.706	28	0.306	0.381			
	27			4.906	1.127	5.398	0.82	28	1.340	0.096			
WAI/GOAL	7			5.188	0.85	4.866	0.871	28	0.394	0.348			
	27			4.708	1.13	5.023	0.966	28	0.617	0.271			
LIKE LAURA	30			4.625	1.598	5.444	1.097	28	2.595	0.007			
CLOSENESS	30			2.25	0.707	2.556	0.784	28	1.478	0.075			
USEFUL LAURA	30			2.25	0.886	2.833	0.924	28	1.774	0.043			
CONTINUE LAURA	30			1.75	0.707	2.611	0.85	28	3.389	0.001			
	44			2	0.926	2.833	1.043	26	1.881	0.036			
FAREWELL	30			0.125	0.354	0.833	0.383	26	4.984	0.000			
MIN/DAY	-6-0	27.53	14.49	17.62	9.091	28.91	28.85	28	1.301	0.102	43	0.177	0.430
	1-7	22.99	9.72	29.21	9.289	30	9.503	28	1.013	0.160	43	0.943	0.175
	8-14	23.64	15.55	32.22	12.82	32.77	15.11	28	0.702	0.244	43	0.399	0.346
	15-21	27.53	15.55	31.11	15.89	37.48	13.77	28	1.493	0.073	42	0.619	0.270
	22-30	24.53	16.27	34.35	11.6	36.34	18.31	28	0.855	0.200	42	1.384	0.087
	38-44	22.27	9.547	29.13	16.88	29.53	19.33	26	0.157	0.438	39	0.503	0.309
DAY/WK>30MIN	-6-0	1.818	2.089	1	1.323	2.471	2.322	28	2.058	0.025	43	0.023	0.491
	1-7	2.545	1.508	3.333	1.5	3.412	1.502	28	0.835	0.205	43	0.095	0.462
	8-14	3	2.49	4.222	1.302	4.176	2.27	28	0.394	0.348	43	1.218	0.115
	15-21	2.818	2.228	4.556	2.007	4.412	1.97	28	0.394	0.348	42	1.920	0.031
	22-30	3.364	2.501	5.556	2.698	5.765	2.635	28	0.297	0.384	42	2.073	0.022
	38-44	2.636	2.248	3.667	2.062	3.294	2.285	26	0.266	0.396	38	0.873	0.194
STEP/DAY	1-7	7438	2376	8940	1883	7763	3298	28	0.977	0.169	43	0.958	0.172
	8-14	7752	2366	9990	2390	8777	4127	28	0.965	0.171	43	1.527	0.067
	15-21	9049	2099	10077	1975	9857	3361	28	0.340	0.368	42	0.845	0.201
	22-30	8011	2766	9984	3490	9795	3815	28	0.134	0.447	42	1.283	0.103
DAY/WK>10KSTEP	1-7	2	1.844	2.111	1.453	2.176	2.099	27	0.499	0.311	42	0.199	0.421
	8-14	1.545	1.809	2.667	1.5	2.353	2.37	27	0.508	0.308	41	1.037	0.153
	15-21	2.364	1.567	3.222	1.563	2.882	2.058	27	0.204	0.420	40	0.451	0.327
	22-30	2.091	1.921	3.667	2.915	3.647	2.827	28	0.389	0.350	41	1.920	0.031
MIN/DAY CHNG	1-30	0.853	3.039	1.433	5.117	2.373	6.248	28	0.366	0.359	42	1.713	0.047
D/WK>30M CHNG	1-30	0.227	0.546	0.7	1.137	0.729	0.895	28	0.225	0.412	42	1.637	0.055
STEP/DAY CHNG	1-30	301.4	980.2	321.7	1040	717.6	1139	28	0.836	0.205	42	0.726	0.236
D/WK>10KS CHNG	1-30	0.109	0.73	0.522	0.88	0.494	0.772	26	0.327	0.373	38	1.738	0.045
DAYS	1-30	19.06	5.272	19.27	5.331	16.63	4.98	28	1.365	0.092	43	0.685	0.248
LOGINS	1-30	24.75	10.34	25.55	7.942	20.89	7.172	28	1.646	0.055	43	0.569	0.286
PAGES	1-30	1.059	0.072	1.133	0.208	1.259	0.283	28	1.294	0.103	43	2.047	0.023
STAGE	Intake	1.923	0.862	1.8	0.919	1.889	0.758	28	0.703	0.244	43	0.610	0.273
	30	2.692	1.109	2.8	0.789	2.667	0.767	28	0.472	0.320	42	0.065	0.474
	44	2.538	1.198	2.7	0.823	2.5	0.786	26	0.635	0.266	38	0.192	0.425
SELF-EFFICACY	1	2.252	0.838	2.061	0.433	2.127	0.448	28	0.129	0.449	43	0.305	0.381
	29	2.53	0.706	2.289	0.574	2.583	0.511	28	1.220	0.116	42	0.002	0.499
	44	2.598	0.892	2.733	0.524	2.71	0.536	26	0.112	0.456	38	0.543	0.295
PROS	1	3	0.673	2.9	0.662	3.233	0.472	28	1.328	0.098	43	0.322	0.374
	29	2.8	0.86	2.74	0.542	3.111	0.697	28	1.221	0.116	42	0.203	0.420
	44	2.8	0.796	2.76	0.44	3.033	0.749	26	1.052	0.151	38	0.403	0.345
CONS	1	3.415	0.69	2.98	0.757	3.344	0.682	28	1.022	0.158	43	1.165	0.125
	29	3.277	0.724	3	0.816	3.356	0.478	28	1.370	0.091	42	0.701	0.243
	44	3.338	0.68	3.04	0.704	3.289	0.63	17	0.930	0.183	38	0.770	0.223
CONTINUE FT	30	3.231	0.599	3	0.816	3.111	0.758	28	0.832	0.206	42	0.420	0.338
	44	2.923	0.862	2.7	0.823	3.056	0.873	26	1.053	0.151	38	0.155	0.439

Table 9-6. Between-group Planned Comparisons for Sedentary Subjects

Measure	Change		ALL CONDS			CONTROL			NON-REL			RELATIONL		
	From	To	df	t	p	df	t	p	df	t	p	df	t	p
	Day1	Day2												
WAI/COMP	7	27	27	0.224	0.824				9	0.940	0.372	17	0.796	0.437
WAI/BOND	7	27	27	0.721	0.477				9	1.332	0.216	17	0.587	0.565
WAI/TASK	7	27	27	0.132	0.896				9	0.605	0.560	17	0.456	0.654
WAI/GOAL	7	27	27	0.275	0.786				9	0.447	0.666	17	0.973	0.344
CONTINUE LAURA	30	44	27	1.867	0.073				9	1.809	0.104	17	1.074	0.298
MIN/DAY	-6-0	22-30	40	1.603	0.117	12	0.232	0.821	9	4.523	0.001	17	0.939	0.361
	1-7	22-30	40	1.910	0.063	12	0.505	0.623	9	1.498	0.168	17	1.311	0.207
	22-30	38-44	40	2.397	0.021	12	0.516	0.615	9	1.487	0.171	17	2.092	0.052
DAY/WK>30MIN	-6-0	22-30	40	6.181	0.000	12	2.185	0.049	9	5.344	0.000	17	4.037	0.001
	1-7	22-30	40	4.348	0.000	12	1.802	0.097	9	2.113	0.064	17	3.424	0.003
	22-30	38-44	40	4.896	0.000	12	1.302	0.217	9	2.683	0.025	17	4.649	0.000
STEP/DAY	1-7	22-30	40	2.529	0.015	12	0.855	0.409	9	1.277	0.234	17	1.999	0.062
DAY/WK>10KSTEP	1-7	22-30	38	2.388	0.022	11	0.137	0.894	8	1.827	0.105	17	1.959	0.067
STAGE	Intake	30	40	6.151	0.000	12	2.993	0.011	9	3.873	0.004	17	3.757	0.002
	30	44	40	1.636	0.110	12	1.000	0.337	9	0.557	0.591	17	1.144	0.269
SELF-EFFICACY	1	29	40	4.819	0.000	12	2.630	0.022	9	1.595	0.145	17	3.823	0.001
	29	44	40	2.582	0.014	12	0.500	0.626	9	4.707	0.001	17	1.129	0.275
PROS	1	29	40	1.872	0.069	12	1.842	0.090	9	0.802	0.443	17	0.878	0.392
	29	44	40	0.364	0.718	12	0.000	1.000	9	0.105	0.918	17	0.548	0.591
CONS	1	29	40	0.437	0.664	12	1.128	0.281	9	0.102	0.921	17	0.098	0.923
	29	44	40	0.000	1.000	12	0.671	0.515	9	0.452	0.662	17	0.595	0.560
CONTINUE FT	30	44	40	1.537	0.132	12	1.075	0.303	9	1.406	0.193	17	0.325	0.749

Table 9-7. Longitudinal Changes for Sedentary Subjects

9.4.1 Relational Measures for Sedentary Subjects

Working Alliance Inventory

Figure 9-14 shows the means for the two administrations of the working alliance inventory questionnaire. The trends are the same as for the full group of subjects (compare with Figure 9-7) but the differences between RELATIONAL and NON-RELATIONAL are even more significant. On the bond dimension at day 7, $t(28)=2.55$ $p=.008$, and at day 27, $t(28)=3.46$, $p=.001$. In addition, at day 27 the overall working alliance composite measure is also significantly greater for those in RELATIONAL compared with NON-RELATIONAL, $t(28)=2.15$, $p<.05$.

Single Item Relational Questions

Subjects reported liking Laura more in RELATIONAL condition, $t(28)=2.60$ $p=.007$ (Figure 9-15), and felt that she was more useful $t(28)=1.77$ $p<.05$ (Figure 9-16) compared with those in NON-RELATIONAL condition. Subjects in RELATIONAL condition also expressed a significantly stronger desire to continue working with her at day 30, $t(28)=3.39$ $p=.001$, and at follow up, $t(26)=1.88$ $p<.05$ (Figure 9-17), relative to subjects in NON-RELATIONAL condition.

Sentimental Farewell

Significantly more subjects in RELATIONAL chose to perform a sentimental farewell compared to those in NON-RELATIONAL, $t(26)=4.98$, $p<.001$ (see Figure 9-18).

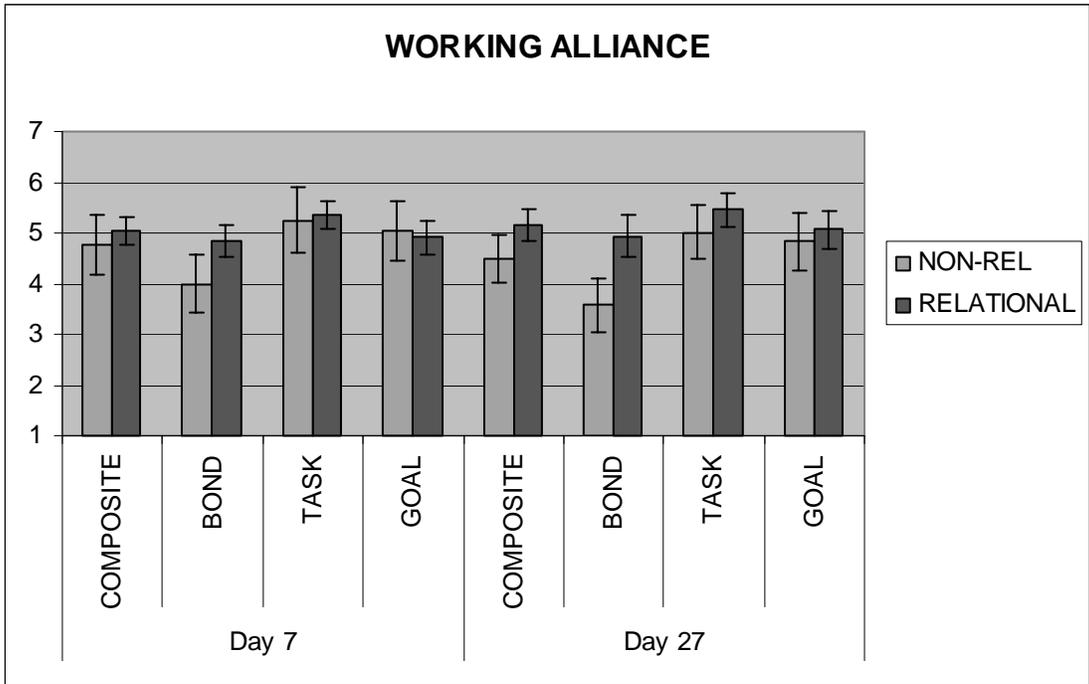


Figure 9-14. Working Alliance Inventory Scores by Sedentary Subjects

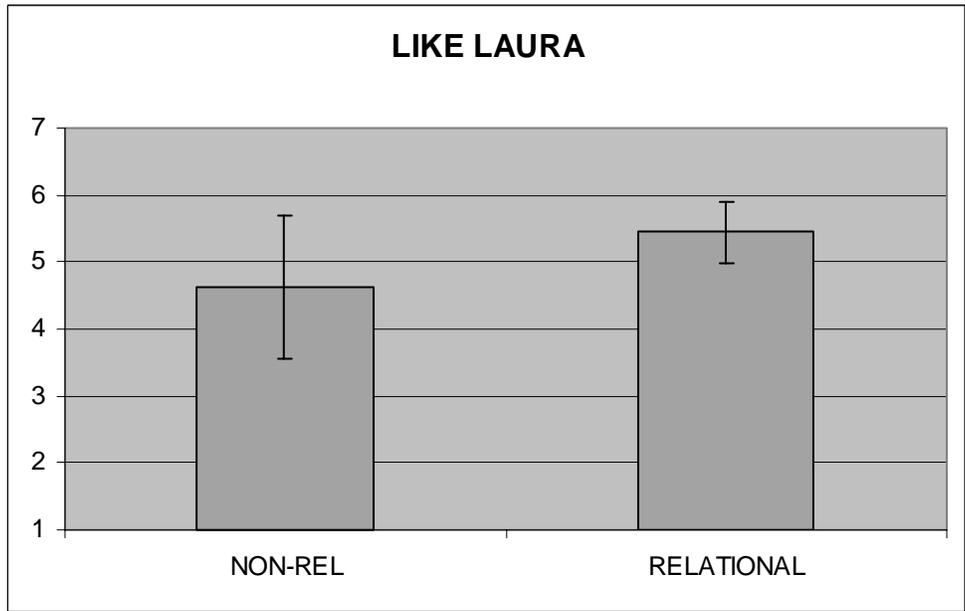


Figure 9-15. Reported LIKING of Laura by Sedentary Subjects

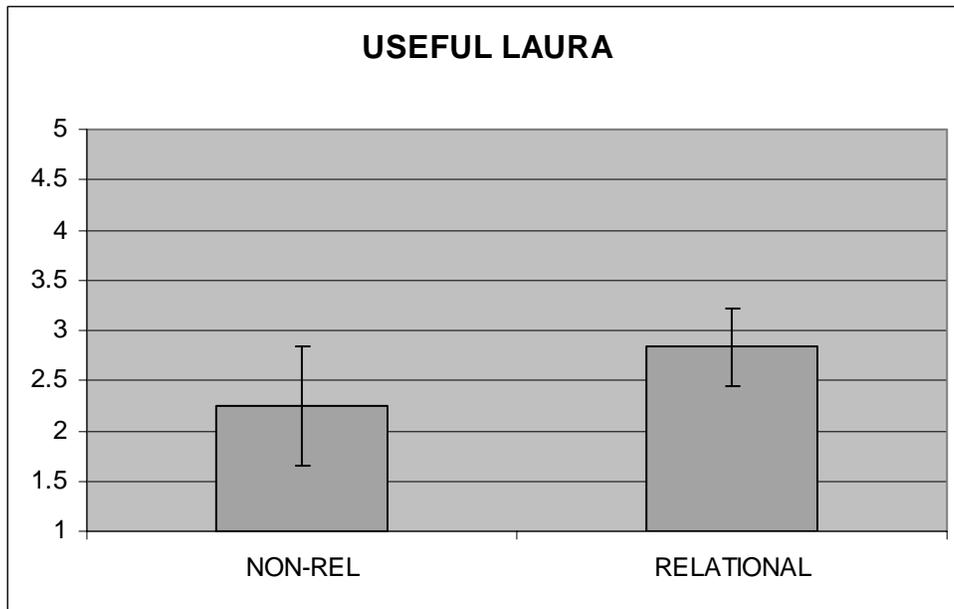


Figure 9-16. Rating of How Useful Laura was by Sedentary

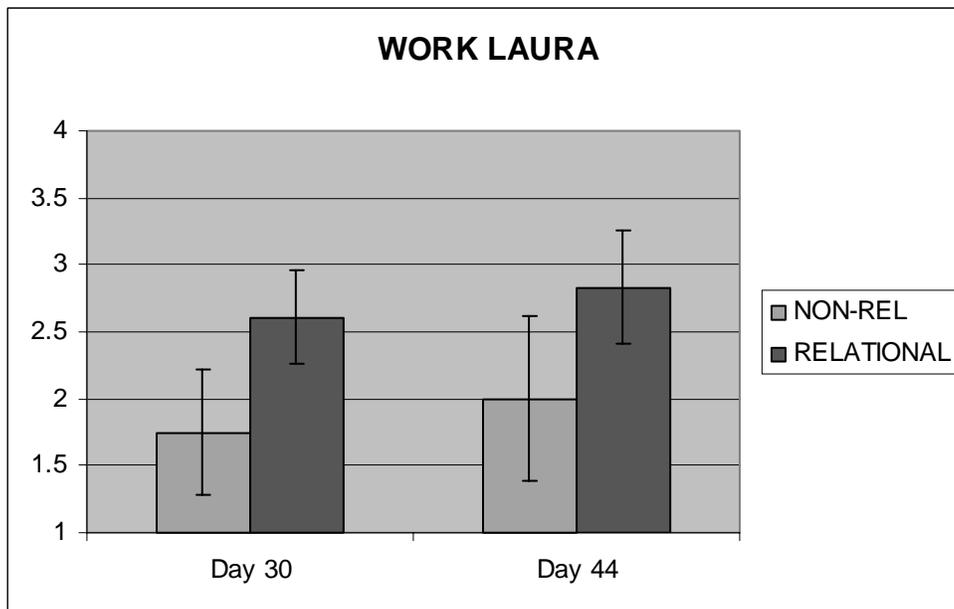


Figure 9-17. Desire to Continue Working with Laura by Sedentary

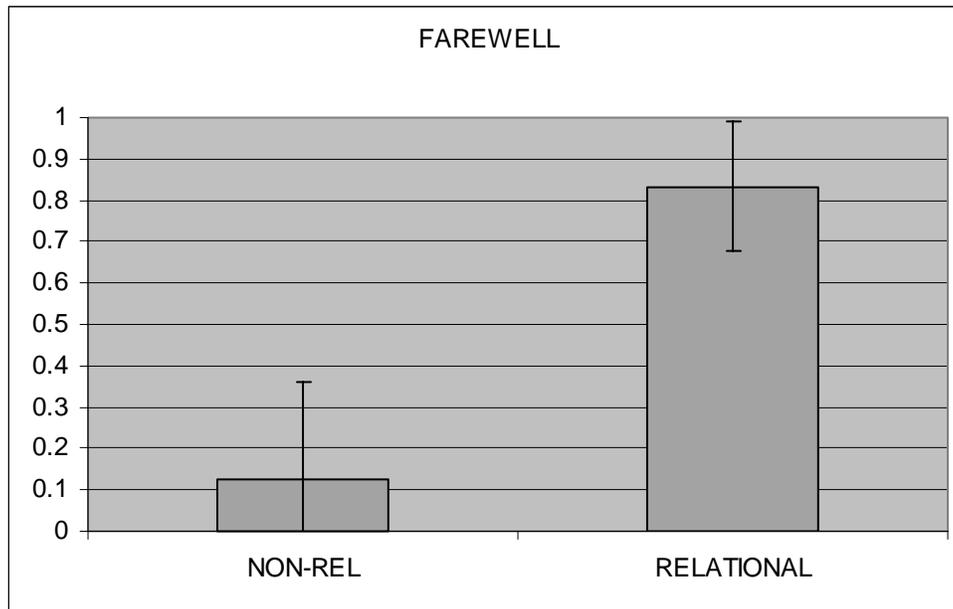


Figure 9-18. Frequency of Sentimental Farewell Use by Sedentary

Helpfulness of Laura

When asked “who was the most helpful in getting you to exercise over the last month?” the distribution of responses was similar to that presented in Figure 9-11, but the differences between votes for Laura, Friends, Family and Workout Buddy were not significant.

Evaluation of Hypothesis H1

Overall, hypothesis H1 received strong support, in most cases even stronger support than for the full set of subjects (presented in section 9.3.1).

9.4.2 Physical Activity Measures for Sedentary Subjects

Minutes per Day of Moderate-or-Better Activity

There were no significant differences between the planned comparison groups for self-reported minutes of moderate-or-better physical activity per day. Subjects in NON-RELATIONAL condition significantly increased their minutes per day of activity during the intervention, paired $t(9)=4.52$ $p=.001$. Across all groups there is a significant decrease in this measure between the end of intervention and follow up, paired $t(40)=2.40$ $p<.05$.

Days per Week Meeting Time Goal (30 minutes/day moderate-or-better activity)

Figure 9-19 shows average days per week that subjects met the target behavior of 30 minutes per day of moderate or better activity. The non-CONTROL groups met the goal level on significantly more days per week than the CONTROL group in week 3 ($t(42)=1.92$, $p<.05$) and week 4 ($t(42)=2.07$, $p<.05$). All groups combined showed significant increases between baseline and the end of intervention, $t(40)=6.18$, $p<.001$, and drop offs between end of intervention and follow up, $t(40)=4.90$, $p<.001$, $d=0.94$, a pattern repeated by most treatment groups (all groups significantly increasing, RELATIONAL and NON-RELATIONAL significantly decreasing).

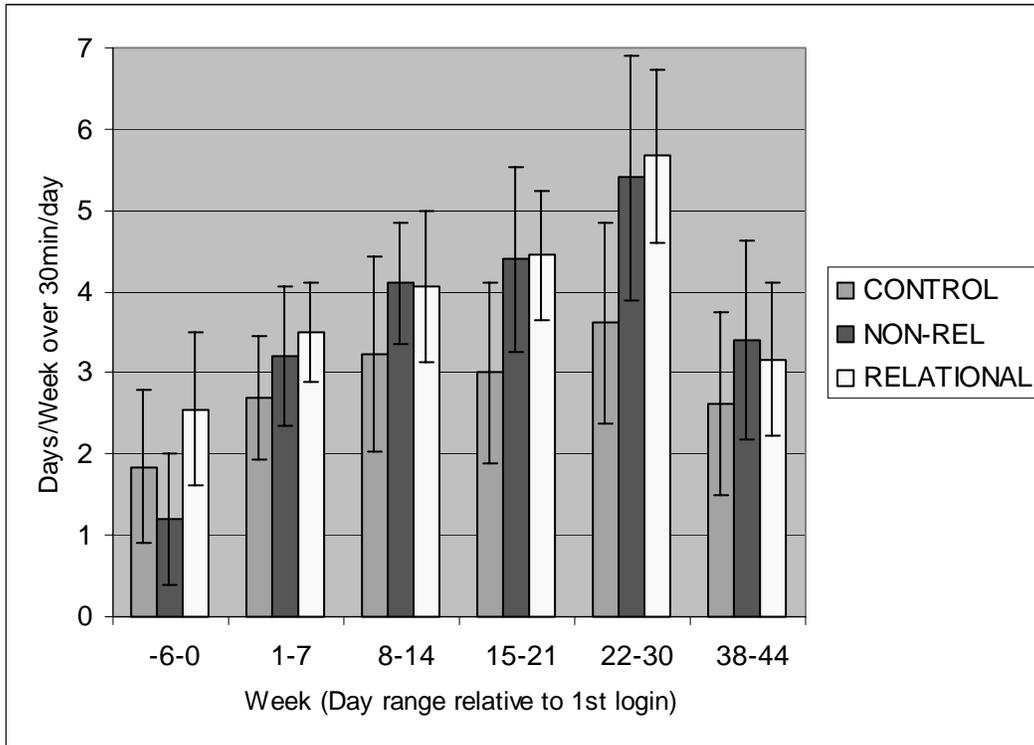


Figure 9-19. Days per Week Over Goal (30 minutes of moderate or better activity)

Steps per Day

There were no significant differences between the planned comparison groups for steps per day measured by the pedometers. There was a significant increase over the four weeks of the intervention for all subjects combined, $t(40)=2.53$ $p<.05$.

Days per Week Meeting Step Goal (10,000 steps/day)

Figure 9-20 shows day per week at or above the target level of 10,000 steps per day. Planned comparisons indicated that the only significant difference between groups was between CONTROL and non-CONTROL groups in week 4, $t(41)=1.92$, $p<.05$. There was a significant increase in this measure during the intervention for all subjects combined, $t(38)=3.00$, $p<.05$.

Rate of Change

Figure 9-21 shows the slopes of least-squares lines fit to each of the four physical activity measures over the four weeks of the intervention (note that slope for steps has been normalized to 1,000 steps per day per week of the study for comparison purposes). Planned comparison analyses indicate that there is a significant difference between the CONTROL group and the other two groups in the rate of change of minutes per day of reported activity, $t(42)=1.71$ $p<.05$ and days per week over 10,000 steps, $t(38)=1.74$ $p<.05$. The difference in rate of change of days per week over 30 minutes of moderate-or-better activity is also

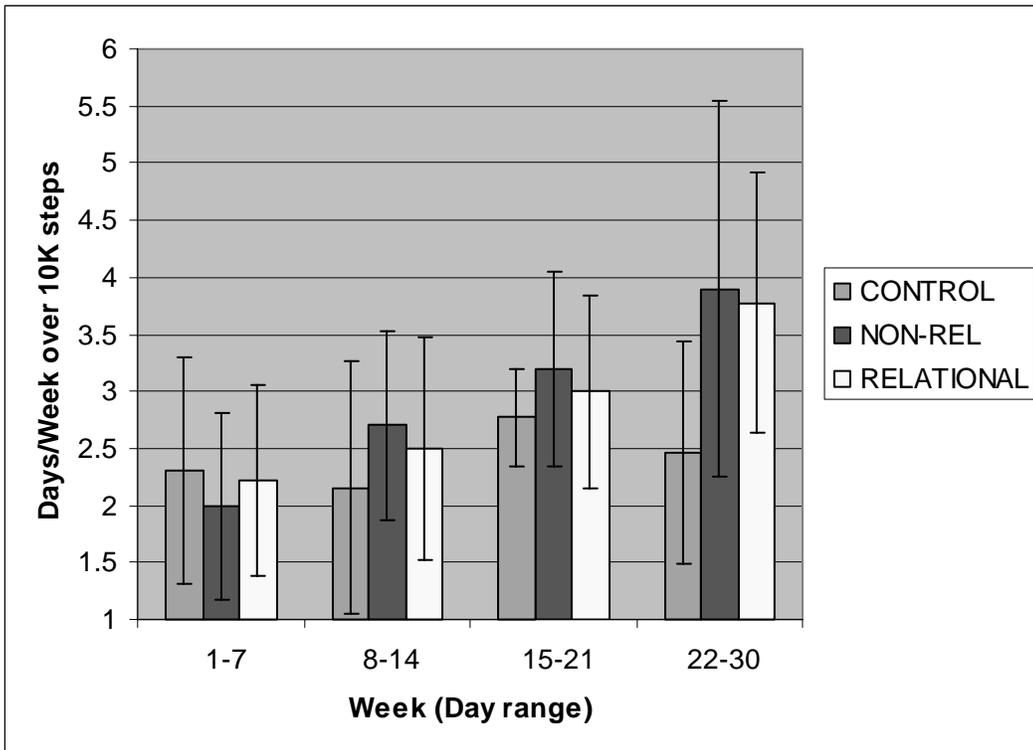


Figure 9-20. Days per Week over 10,000 Steps

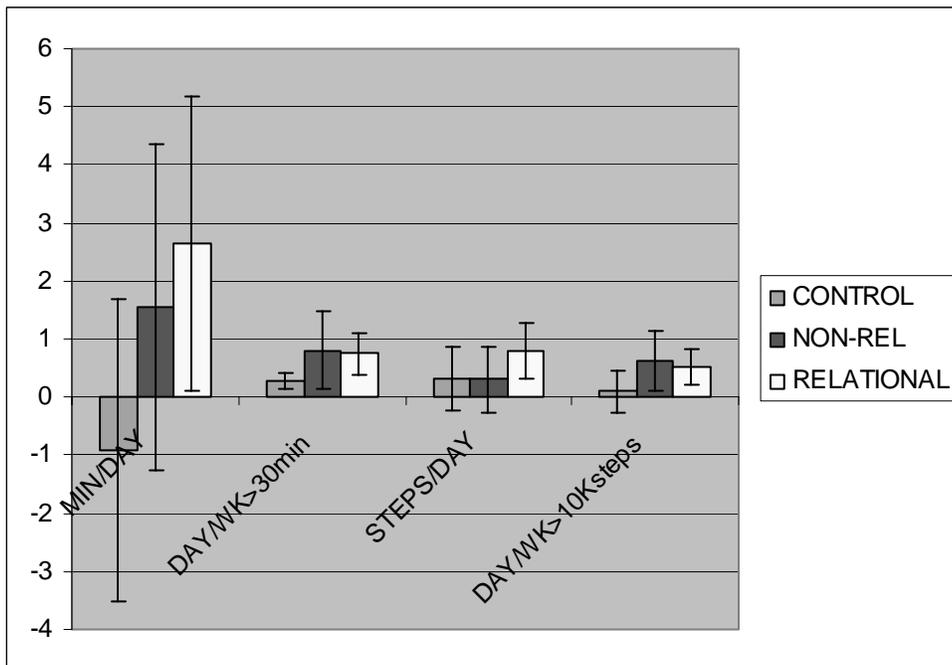


Figure 9-21. Rate of Behavior Change per Week of Intervention

approaching significance for the comparison between the CONTROL and non-CONTROL groups, $t(42)=1.64$ $p=.055$.

Evaluation of Hypothesis H2

There was no support of hypothesis H2 (a), that subjects in RELATIONAL condition would outperform those in NON-RELATIONAL condition. However, the data suggest that there is some support for H2(b); groups with the agent outperformed the CONTROL group at two time points on the measure of days per week over 30 minutes of moderate activity, and at one time point on the measure of days per week over 10,000 steps, and on two measures of rate of behavior change during the intervention period.

9.4.3 Participation Measures for Sedentary Subjects

Attrition

Three subjects from the sedentary analysis group voluntarily dropped out of the study, two in CONTROL condition and one in NON-RELATIONAL condition.

Logins

There were no significant differences between treatment groups with respect to either number of days logged in, or total number of logins.

Pages Viewed per Session

Figure 9-22 shows the number of educational content pages viewed per session by treatment group. The difference between CONTROL and non-CONTROL groups is significant, $t(43)=2.05$, $p<.05$.

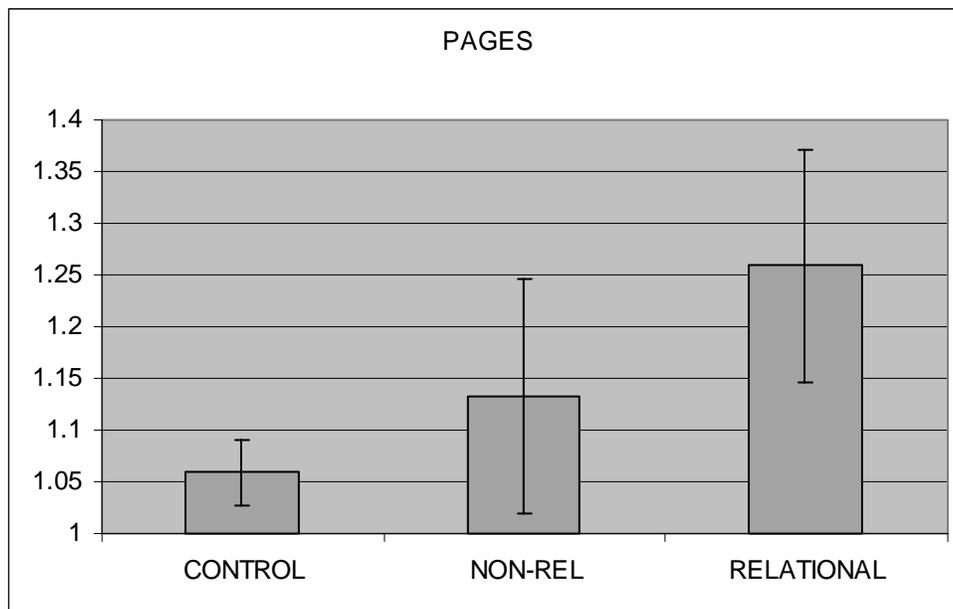


Figure 9-22. Educational Content Pages Viewed per

Evaluation of Hypothesis H3

As in the prior analysis there is no support for H3 (a); subjects in RELATIONAL and NON-RELATIONAL groups participated at approximately the same levels in the FitTrack program. The data suggests some support of hypothesis H3(b), given that subjects in non-CONTROL groups viewed more educational content pages than those in the CONTROL group.

9.4.4 Exercise Attitudinal Measures for Sedentary Subjects

Exercise Stage of Change

Exercise stage of change followed the same general pattern as for the full group of subjects. There are no significant differences between planned comparison groups. As before, there is a significant increase in stage of change from intake to the end of intervention, $t(40)=6.15$, $p<.001$ for all subjects combined, as well as significant increases for each treatment group independently. There was no significant change in stage between the end of intervention and follow up.

Exercise Self-Efficacy

There were no significant differences between planned comparison groups. As before, self-efficacy is increasing for all subjects between the start and end of intervention, $t(40)=4.82$, $p<.001$, and between the end of intervention and follow up, $t(40)=2.58$, $p<.05$. Subjects in CONTROL and RELATIONAL groups also significantly increased during the intervention period, while subjects in NON-RELATIONAL increased significantly during the follow up period.

Exercise Decisional Balance

There were no significant differences between planned comparison groups for either the PROs or CONs of exercise. There were also no significant changes over time for this measure.

Evaluation of Hypothesis H4

There was no support for hypothesis H4 (a) or H4 (b); neither the presence of an agent nor the use of relational strategies had an impact on exercise attitudes for the sedentary group of subjects.

9.4.5 Other Quantitative Measures for Sedentary Subjects

System Evaluation

Subjects in this analysis were generally more positive about the FitTrack system relative to the full group of subjects discussed in section 9.3, especially in their rating of how useful Laura was. Table 9-8 lists their ratings of usefulness of different system features. There were no significant differences between groups on any of these measures with respect to the planned comparisons.

	Educational Content	Monitoring Charts	Laura
Extremely	8%	25%	3%
Very	25%	33%	12%
Moderately	27%	25%	40%
Slightly	40%	13%	36%
Not at all	0%	4%	9%

Table 9-8 Usefulness of FitTrack System Features

9.5 Qualitative Analyses

Interviews were held with 28 subjects (16 in RELATIONAL and 13 in NON-RELATIONAL) during the follow up period or just before final debriefing to get a qualitative sense of their reaction to the FitTrack program and Laura. In addition, at the end of the intervention period and again at follow up, subjects were presented with an open-ended feedback form that they could use to provide comments and suggestions, and many subjects took the time to provide substantive input (28 in CONTROL, 21 in NON-RELATIONAL and 29 in RELATIONAL).

9.5.1 Overall Impressions

Aside from issues with Laura and problems with particular features, the overall impression of the FitTrack program was very positive. Some representative responses when asking subjects what they thought of the overall concept were:

I guess I thought...it's a really really good idea, and some aspects of it were very helpful, like being able to record your progress and things like that. (RELATIONAL)

It was useful. It was my number one motivation for working out. It was a consistent motivation for working out. (NON-RELATIONAL)

It was really good. I was exercising much more than before. (RELATIONAL)

Negative responses tended to be from people who felt that the program wasn't appropriate for them, because they were already engaged in a sport or exercise activity that the program didn't cover ("As for me, I like swimming, but the selections don't have the answer I want to choose.") or because they felt that it wasn't otherwise tailored to their particular lifestyle ("It was very definitely designed for college students, with all this talk about getting out of your dorm room and exercising with a buddy along the emerald necklace. None of that applies to my life.").

The most frequently mentioned feature of FitTrack (aside from complaints or compliments about Laura) was the usefulness of wearing pedometers and tracking progress via the self-monitoring charts to keep track of actual activity:

I think wearing the pedometer was a huge thing. It's a number, so it was useful, it was very visual feedback about how much I'd been exercising. (RELATIONAL)

I liked the progress charts, I thought that was good. because it helped me to see what was happening over the course of a few weeks. (NON-RELATIONAL)

Three subjects (all in RELATIONAL condition) even pleaded with the experimenter to keep the program running so that they could use it after the study had ended:

My only other comment is, is this program going to be available to study participants? Can I keep working with it? Because I know that sometimes with studies you are given the benefit of using the program. Like, you guys are letting us keep the pedometers, and I want to continue... I'm still keeping track of my steps most days, but I would keep track every day if I had the program still. Whereas now, it's about every other day that I remember to write down the number. [Interviewer: I can look into keeping it running.] That would be really cool. (RELATIONAL)

9.5.2 Acceptance of Laura

Many subjects reported that the interactions with Laura took a little getting used to. Three subjects used the exact same phrase in describing it: "It was a little weird at first, but I got used to it." However, after a short time, subjects seemed to adapt to Laura, and vice-versa:

At first I don't know if I took the character that seriously. Over time, I was like, OK this character clearly remembers things. (RELATIONAL)

It was sort of artificial at first. It was very generic. But after awhile she adapted to me, which was good. (NON-RELATIONAL)

When asked if they felt uncomfortable in their interactions with Laura, not a single subject responded affirmatively, although some qualified their remarks by saying that, while they were never uncomfortable, their interactions were sometimes "frustrating", "bizarre" or "weird".

One subject also mentioned that seeing their behavior change over time made them trust Laura more:

I didn't feel like I knew her initially. But after you get to know her a bit it got better. But when you see your charts and see that you are actually improving over time, that makes you trust her. (NON-RELATIONAL)

Regarding whether they liked the overall concept of conversing with and relating to an animated character, subjects reported strong opinions on both sides of the issue. Some of the representative positive responses included:

It was a really, really great idea to have some kind of animated character because it makes you feel like you're actually talking to a person rather than having words on the computer screen. (NON-RELATIONAL)

I like talking to Laura, especially those little conversations about school, weather, interests, etc. She's very caring. Toward the end, I found myself looking forward to these fresh chats that pop up every now and then. They make Laura so much more like a real person. (RELATIONAL)

I wish she could imitate a real person's life in her answers rather than sticking to the reality and saying things like she is limited to that box. Maybe this has something to do with trainees wanting to have role model to achieve their own physical fitness roles by taking the trainer as a role model. Or maybe it is just about having a richer conversation helping getting connected to the other person. (RELATIONAL)

Once I kind of got used to Laura in general, I didn't really see her as a computer character. It didn't really bother me. (NON-RELATIONAL)

Some subjects clearly did not like the idea of interacting with a computer character:

It was hard to divorce the notion that Laura was a computer simulation from my interactions with her, so it was a little weird to try to be "personable" with her. (RELATIONAL)

I didn't really like Laura very much. ... Actually, I liked all of the software except for the animated conversation thing. (RELATIONAL)

Personally I detested Laura. (NON-RELATIONAL)

Laura is NOT a real person, and therefore I HAVE NO RELATIONSHIP WHATSOEVER WITH HER! (NON-RELATIONAL)

The use of multiple choice menus for input seemed to a few users to unnaturally limit what they could say in the conversation, and how they could say it:

The software is fine, but I had some problems with Laura because she didn't quite understand what I was doing. There weren't enough choices, in the questions I could ask her... Like when I got sick, I didn't know how to explain that to her... That's why I would prefer to talk to a person rather than a robot. (RELATIONAL)

Although, there were definitely times where I wanted to say something and it wasn't one of the choices. (RELATIONAL)

Finally, while several users complained about Laura's robotic synthetic voice ("I thought her voice was a big turn off."), some actually said that it seemed appropriate and empathetic ("She had a very kind voice that was very encouraging, that was a good feature.").

9.5.3 Repetitiveness

The most frequent complaint about Laura, and the system in general, was the repetitiveness of the dialogue. Even though there were almost 1,000 states in the dialogue state machine, each of which typically had two or more responses that could be selected at random, most subjects still felt that, at some point during the study, the interactions became repetitive and predictable for them. This problem was especially acute for subjects in the NON-RELATIONAL group, since there was a much narrower range of dialogue that they engaged in, and it did not vary much based on historical context (relative to the dialogue in the RELATIONAL condition).

The first couple of days I was impressed by it. But, there didn't seem to be a lot of variety going on after that, so it kind of lost my interest, it lost the engagement factor. Maybe, six or seven days into the study I could almost predict what she was going to say, and once the engagement was lost you sort of lose the power of the animated instructor. ... (NON-RELATIONAL)

Like 15 days into the study when I could almost predict what she was going to say, it became easier to do things like check my mail in between her responses. ... Even with just little bits of variety your mind doesn't shut off. (NON-RELATIONAL)

After awhile I'd click on random things just to see what she would say. Otherwise she would say the same thing, so there was no sense clicking on it. (NON-RELATIONAL)

Sometimes I would just get a bit frustrated because she was asking the same things every time in the same order, it was just more difficult to stay focused on what I was doing there. (RELATIONAL)

Laura is waaaaaaayyyyyy too repetitive. (RELATIONAL)

In addition to Laura's responses seeming repetitive, some subjects also complained about lack of variability in the dialogue moves they could make. Overall, Laura's repetitiveness did more than simply annoy subjects; several mentioned it as significantly decreasing the ability of the system to motivate them to exercise:

In the beginning I was extremely motivated to do whatever Laura asked of me, because I thought that every response was a new response. Whereas, towards the end I could tell what she was going to say to a couple of my responses. (RELATIONAL)

The negative aspects of it were that Laura was very repetitive, so it was actually more motivating in the beginning to talk to her than later on, which is sort of the opposite, I think, of what is intended. Because she would go through the same routine every single time, so it wasn't very realistic. As a result I didn't feel obligated, I didn't feel like I had to impress her in any way. (RELATIONAL)

9.5.4 Caring

One of the goals of building a working alliance is for a client to establish the perception of the therapist as genuinely caring about them. When asked whether they felt that Laura cared about them, many subjects responded affirmatively but qualified their responses to acknowledge that Laura wasn't a real person, and thus couldn't genuinely care for them:

Yes, as much as a computer can care. (RELATIONAL)

Yea, I think there was an illusion there that she did. (RELATIONAL)

As much as it mattered to ... I never forgot that it was a computer program, but you'll notice that I find myself calling her by feminine pronouns rather than calling her an 'it'. So, I definitely remembered that she was a computer program, but I did feel like it was a more personal interaction than that. (RELATIONAL)

Other subjects were even less sure about whether a computer could "have" feelings such as caring:

She's a computer character. I don't know if she cared about me. I don't know if she feels. She's a character and has a role, but I don't know if she has feelings. But, it worked for me and I'm happy. (RELATIONAL)

I find 'care' to be a funny term to use with a computer character. I felt like it was helpful to have positive reinforcement, even if it was from a computer character. (NON-RELATIONAL)

And, of course, there was a group of subjects who put feelings such as caring firmly outside the realm of possibility for a computer:

No, not really. I think it's because I knew it was a computer. (NON-RELATIONAL)

No. I felt like I was talking to a robot, to a machine. (RELATIONAL)

No, not really, because I plugged in a number and she had a script. (NON-RELATIONAL)

It's hard to say, because I'm dealing with computers all the time. So, I really cannot take her as a character when you talk about those emotional kinds of things. So, not really. (RELATIONAL)

One subject even interpreted Laura's caring as a function of the emotional investment of her developers; "I could see the care the programmers put into her."

Finally, one subject made a connection between their perception of Laura's repetitiveness and their assessment of her caring for them:

...at the beginning her expressions and tone of voice were somewhat realistic, but then just like if I were trying to motivate someone and said the exact same thing every single time, they might get the impression that I cared about them early on, but later on, if I'm saying the exact same words, then at some point they are thinking that I don't really care about them. (RELATIONAL)

9.5.5 Laura vs. Human Trainer

When asked whether they would have rather interacted with a human trainer than with Laura, subjects gave a wide range of opinions. Most subjects cited convenience as the primary reason why they would prefer Laura:

No, not given my time constraints. With respect to a webcam and a live trainer, that would be a toss up, but I don't think you could get one to be available at 11pm. (NON-RELATIONAL)

No, I liked the convenience of being able to meet with her in my own time, when I was free and ready to sit down and do it. It would have been harder to match schedules with a person. (RELATIONAL)

That would have been useful, but I could talk to Laura anytime. (RELATIONAL)

Some also said they would prefer Laura over a human trainer because her responses were predictable, especially when they hadn't done well on their exercise:

It was kind of easier to interact with the computer character, just because depending on how I did I knew what her responses would be. So, I guess it was easier to interact with a computer character. (NON-RELATIONAL)

Some subjects also indicated that they would prefer interacting with Laura to interactions with a personal trainer because they feel less guilty about letting her down if they were not able to exercise:

Probably not. More painful that way, if I decide not to do something. I've interacted with live people before and given up on them because of exercise. Because sometimes I want to go off the wagon for awhile on exercise. And, when talking with a computer program about exercise, I never want to avoid that person in real life if I haven't made my exercise goal. So, I kind of liked the fact that she was a computer program and not a person. (RELATIONAL)

One subject also indicated that they thought the combination of Laura and a personal trainer would be ideal:

this worked for me amazingly well. ... With this interaction however, I started to think that a computer character can accomplish a large amount of daily responsibilities that a human personal trainer can do. I would be happy to have first a computer character work with me and then I could be referred to a human trainer for my 'other' type of questions (that can't be answered by selection of phrases provided to me). (RELATIONAL)

The reasons given for preferring a human trainer over Laura ranged from Laura's inability to understand subjects, her repetitiveness, her inability to go with them when they exercise and know whether they are paying attention or not, and the perception that they would be more motivated by guilt to exercise for a human trainer:

I would have. This way I would be able to explain everything that I wanted to. (RELATIONAL)

She had some good points when I said I couldn't get out to exercise, but it felt to me that there were extenuating circumstances from time to time, I work at home and I have two small kids, and it makes it more difficult. That's the kind of things I could explain to a human trainer and there might be other ways around it, but Laura wasn't flexible enough for me to be able to explain my circumstances. (RELATIONAL)

Yes, but simply because of the repetitiveness issue. (RELATIONAL)

Often a trainer is with you, so if Laura was on a PDA with me when I was exercising, that's a total different dimension that a personal trainer can give you, that she couldn't right now. (RELATIONAL)

A trainer would know if I was paying attention or not. (NON-RELATIONAL)

It would have given me more incentive. It would have been more real. ... I would have felt worse about letting them down. Since it was an animated character that didn't have feelings it didn't matter so much. (NON-RELATIONAL)

9.5.6 Laura as Motivator

Several subjects talked about Laura's ability to motivate them. Most said that they felt responsible to her for meeting their goals, and would feel guilty if they hadn't met them:

Because I knew I had to enter the numbers every day, it was like a responsibility to someone else. (RELATIONAL)

When I said I couldn't exercise I felt bad. When she said "are you sure you can't exercise?" it would make me think about it. (NON-RELATIONAL)

It sort of kept me motivated, because I always do more if I know I'm responsible to someone. (RELATIONAL)

It kept you on your toes because you didn't know if you were going to meet with the animated person. (RELATIONAL)

As silly as it sounds, I find that I found a little motivation to exercise knowing that Laura would ask if I did or not. Now that I don't have anyone checking, I find it harder to get motivated. (RELATIONAL)

It is motivation to know that you'll have to tell Laura you didn't exercise like you planned. (NON-RELATIONAL)

There were a few subjects who expressed the opposite point of view ("I didn't feel like, I'm going to let the computer program down if I don't exercise."), and some who saw Laura's inability to make them feel guilty as a plus, because they wouldn't have any hesitations interacting with her and could be completely truthful with her about not meeting their exercise goals:

I did not have to worry about the impression that I will make on the personal trainer, how the other person is actually doing that day, she will always be there and supportive guaranteed, she would not be bored with me or disappointed with me sometimes not attaining my personal exercise goals. (RELATIONAL)

It worked a lot better than... like most of the time in my past it's been my mom saying "you need to exercise", or my friends saying "why don't you come lift weights with us". Whereas, with the

computer program I knew that if I never wanted to have repercussions for not exercising, I could tell her truthfully how much I've exercised, because at the end of the day she wasn't going to come out of the computer and berate me for half an hour for not exercising enough. So, I felt like I could very easily, truthfully say how much I've exercised. (RELATIONAL)

Better than someone else telling me. My mom or dad telling me to exercise is embarrassing, this is much easier. (RELATIONAL)

Laura's ability to make subjects feel guilty about not exercising led some of them to alter their behavior so that they would delay interacting with her until they had met their goals:

And, I also figured out that if I hadn't made my goal for the day, if I just waited until tomorrow to log in and I made the goal tomorrow, she wouldn't care. In a way that gave me an easy out, but in a way that was kind of disappointing, because she didn't ride my ass if I did like half of my goal the previous day she made no comment whatsoever. So, I felt like I was getting away with something. (RELATIONAL)

9.6 Discussion

Across both analysis groups (all subjects and sedentary subjects) there was strong support for hypothesis H1; the use of relational strategies did result in significant increases in a number of measures of the quality of the relationship between subjects and the agent. However, there was little or no support for most other hypotheses, although the data did suggest that the use of an animated exercise advisor who worked with sedentary subjects to set and follow up on daily exercise goals had a significant impact on their physical behavior.

9.6.1 Effects on Relationship

Hypothesis H1 was the most important with respect to the goals of this thesis. Demonstrating that an agent can build and maintain a relationship with a user is a pre-requisite for demonstrating that such a relationship can impact task outcomes, and this study has accomplished that. In analyses involving both all subjects and the target subset of subjects, relational strategies were shown to have significant impacts on the bond dimension of the working alliance (measured at two times during the intervention), on subjects' liking of Laura, and on their desire to continue working with her at the end of the intervention period. The fact that there were no significant changes in these measures over time indicates that Laura was able to maintain the sense of relationship with subjects over the entire month of the intervention.

The most consistently significant items in the bond dimension of the working alliance for both groups of subjects were: "I believe Laura likes me.", "I believe Laura is genuinely concerned about my welfare.", "Laura and I respect each other.", "I feel that Laura appreciates me.", "Laura and I trust one another.", "My relationship with Laura is very important to me.", and "I feel Laura cares about me even when I do things that she does not approve of." These indicate that Laura was able to build a caring, trusting relationship with users when she used the appropriate relational strategies.

Establishing the bond dimension of the working alliance is a pre-requisite for establishing the task and goal dimensions of the alliance; trust must be in place before a client will self-disclose enough to genuinely negotiate therapeutic goals and tasks. Very little work was done on implementing dialogue to affect the task and goal aspects of the alliance, however. The dialogue required to increase these alliance dimensions involves extensive negotiation of

therapeutic goals and tasks, as well as a great deal of flexibility in what the agent can handle in terms of the range and level of customization allowed in the goals and tasks of therapy. Laura pretended to negotiate a little, but in the end the goal was always the same: 30 minutes of activity and 10,000 steps per day. In addition, Laura’s intervention strategies (as well as the educational content) were all designed around these goals, which may be why most of the subjects who had baselines over 30 minutes/day and were given the option of increasing their activity goals on day 8 did not take advantage of this feature.

Relative to prior studies on the effect of working alliance on outcomes in psychotherapy, these significant increases in alliance scores did not translate into expected changes in attitudes or behavior. Perhaps the most relevant outcome measures from prior studies are treatment participation and premature termination—which translate into attrition rates in the FitTrack study—and compliance with medication—which translates into days per week at the recommended goal levels for physical activity. No significant changes in either of these measures were demonstrated to be correlated with changes in working alliance in the present study. One reason for this may be the lack of significant between-group differences in the task and goal dimensions of the alliance, but there are many other possible reasons for the lack of between-group differences in outcome measures, discussed further in the next section.

In addition to the effects on trust and working alliance, relational behaviors may have contributed to behavior change in other ways as well. Talking about off-task topics (small talk) allows the agent another means for keeping variability in the interactions high, thus maintaining subjects’ engagement in the behavior change over time. Having the agent remember things about subjects also allowed the dialogue to become more tailored to their personal lifestyle, for example suggesting that they reward themselves for a week of meeting exercise goals by going to a movie (if they liked movies) or going out for Chinese food (if they liked Chinese food).

9.6.2 Effects on Physical Activity

Even though there were very few significant between-group differences found, all groups in the sedentary analysis increased their levels of physical activity (on several measures) during the intervention period, and decreased their levels of activity immediately following removal of the intervention. Some of the reasons for the increased activity may be simply that being in a study or having to report in every day on behavior are motivating in and of themselves. However, even the CONTROL group used many “state-of-the-art” intervention techniques, including self-monitoring, goal-setting, and education, thus the increase in activity (as well as self-efficacy) for all groups during the intervention period most likely reflects the effectiveness of these proven techniques.

The significant drop off in exercise behavior during the brief follow up period indicates that a lasting change in behavior had not been achieved. Sixty-three percent of subjects who completed the study reported levels of activity at or below their baseline levels at follow up. The drop off was especially acute for those in RELATIONAL condition, as Figure 9-23 shows (note that the between-group differences are not statistically significant). According to one expert in health behavior change, a rapid increase in behavior change during intervention followed by a rapid decrease following the removal of intervention is characteristic of face-to-face interactions with behavior change professionals (Prochaska, 2003). By this measure, it would seem that the RELATIONAL agent has succeeded in replicating some of the effects of

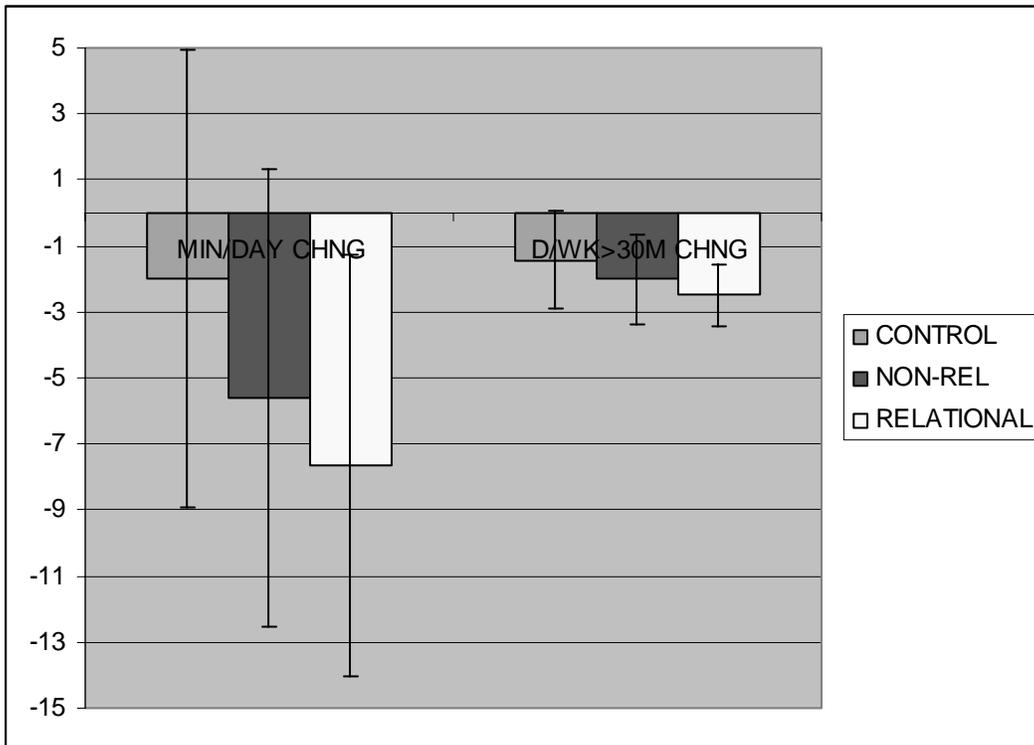


Figure 9-23. Behavior Change between End of Intervention And Follow Up for Sedentary Subjects

face-to-face counseling. One way to reduce the rapid relapse rate is to gradually “wean” subjects off the counselor by having them reduce the frequency of their interactions before terminating the intervention completely.

Another explanation for the rapid relapse observed is the timing of the study—the follow up period coincided with the end of semester and the onset of winter weather in Boston, and for many subjects the final 7-day recall follow up measure spanned the Thanksgiving holiday. Nevertheless, the ability for this kind of program to produce lasting change has yet to be demonstrated. However, an argument can be made that computerized interventions, such as FitTrack, never need to be removed from subjects’ lives. Compared with a human trainer or exercise advisor, Laura could be a permanent part of peoples’ lives, always available, always reliable, with only minimal on-going costs to maintain the relationship.

The lack of significant differences between RELATIONAL and NON-RELATIONAL groups with respect to physical activity measures may be due to several factors. As mentioned above, the establishment of a working alliance bond may be insufficient in and of itself to translate into changes in attitudes and behavior; significant changes in the task and goal dimensions may also be required. It may simply be a matter of too few subjects; the initial power analysis indicated that 60 (30 per condition) would be required, while only 41 subjects in the sedentary group actually completed the study. This is also supported by the fact that 7 of 8 planned comparisons of the rates of behavior change for the sedentary group were in the hypothesized direction, but only 2 of them were statistically significant. Finally, the study was likely too short in duration to detect any real long-term changes in exercise behavior. Attrition is probably one of the most important measures of success (or lack thereof) in this kind of

program, and a study with a much larger set of subjects over a much longer period of time would be required to detect significant changes in this metric. The fact that there were significant increases in self-efficacy for all groups during the intervention also indicates that later changes in behavior are likely, as changes in attitude usually precede changes in behavior by varying lengths of time.

Overall, there were few differences between those who did best in the study (with respect to physical activity) and the average subject. An analysis of the 27 subjects who performed the best (were in the top 80% of all subjects in increase of steps/day or minutes/day of activity during the intervention) indicated that they were typical with respect to all demographic and personality measures. However, there were two measures that seemed to differentiate these performers from the rest of the subjects. First, they were in a significantly later stage of change: 19% were in contemplation, 48% in preparation, and 30% in action vs. the average 29%, 31%, and 25% breakdown, respectively. Second, they seemed to be able to rely on their friends more to help them exercise: 19% of the top performers said that their friends were the most helpful in getting them to exercise while only 11% said that Laura was the most helpful; compared with 11% and 20% ratings by the average subject, respectively.

9.6.3 Repetitiveness

The fact that most subjects found the dialogue repetitive by the end of the intervention, despite significant variations authored into the dialogue scripts, is a very interesting phenomenon and one that warrants further research. The perceived repetitiveness was not just an annoyance; several subjects said that it had a significant negative impact on their engagement with the system and motivation to exercise. Repetitiveness may also have played a role in the effects of relational strategies, since the RELATIONAL condition of the study contained significantly more variation in dialogue than the NON-RELATIONAL condition.

Quantification of the amount and type of variability in verbal and nonverbal behavior required to avoid perceived repetitiveness represents an interesting and important area for future research. This is an especially relevant problem for the field of natural language generation, since generation methods are ultimately required to achieve the broadest range of variability in surface utterance forms. Variability is also important in the input options available to the user.

Nonverbal and paraverbal behavior may have played a role in subjects' perception of repetitiveness, since Laura had a fairly small repertoire of nonverbal behaviors at her disposal and, with the exception of idle behaviors, no variation in the surface form of these behaviors (e.g., she only had one contrast gesture, one left deictic gesture, one gaze away animation, etc.). And, while there was some intonation control available on the speech synthesizer used for Laura's voice, many crucial controls were not available, such as pitch contour and variability. There was slightly more nonverbal behavior variability in the RELATIONAL condition, given that immediacy behaviors translated into more frequent gestures and facial animation, and that the EMPATHY frame was not experienced by subjects in the NON-RELATIONAL condition. However, with the exception of the "concerned" facial expression in the EMPATHY frame, the range of possible nonverbal behaviors was identical across the two conditions.

All types of repetitiveness are not necessarily bad. For example, repeating certain types of therapeutic messages to a client until the client begins to internalize them is an important

technique in cognitive behavioral therapy. However, the determination of which kinds of repetitiveness are beneficial and which kinds are detrimental, with respect to some outcome measure, remains an important empirical question.

There are several possible approaches to avoiding repetitiveness in systems that users must interact with over a long period of time. The first is simply to increase the space of possible dialogue moves. In the approach taken in FitTrack, this means increasing both the size of the dialogue networks and the number of utterances available at each node (representing variations on a particular output message). The development and testing of such an extensive dialogue space represents a very large, but important, task when building a system capable of many interactions with a user. More generally, a text generator that had a variety of pragmatic parameters could be used to vary the agent's output based on subtle changes in context, such as the weather, the subject's performance, or the agent's "mood" (randomly generated each day). Finally, the agent could also apologize and joke about its repetitiveness as a way of alleviating negative attitudes towards it.

9.6.4. Reactions to Laura and FitTrack

As anticipated, reactions to the idea of conversing with and relating to an animated character ranged widely, from completely embracing the idea to abhorring it. However, no subjects reported feeling uncomfortable in their interactions with Laura. More importantly, the experiment provided significant support to the notion that anthropomorphic interfaces are intuitive and natural. During intake, subjects were simply told that there would be an animated character in the software but they were not given any instructions on how to interact with it. The fact that all 100 subjects were able to readily and regularly conduct conversations with this agent without ever requesting help or guidance from the experimenters says something about the intuitive appeal of these kinds of interfaces.

Although previous studies have found significant variations in reactions to anthropomorphic interfaces based on subjects' personality (e.g., Chapter 5), there were almost no significant impacts of subject personality on any of the measures in the study. The only relationship approaching statistical significance was that submissive subjects tended to pursue social dialogue more frequently than dominant subjects when given the choice of opting out. However, this could have simply been a function of their perceiving there to be a social demand in these situations to do social dialogue (opting out was impolite) and submissive subjects simply decided to go along with it more frequently. Regarding the relationship between introverts and extroverts described in Chapter 5, the same general pattern of relationship was found in this study (Figure 9-21 shows this relationship for the bond dimension of the working alliance evaluated on day 7 of the intervention), however, it is far from statistically significant.

The overall reaction to the FitTrack program was very positive. When subjects in the target group were asked (at the end of the intervention period) if they would like to continue using the system, only 2% said "not at all" and 29% said "very much", with the rest falling in between. During interviews, several subjects said that the program had been instrumental in getting them to exercise, as reflected in both behavioral outcome measures and in attitudinal measures (both exercise stage of change and self-efficacy significantly increased over the intervention period).

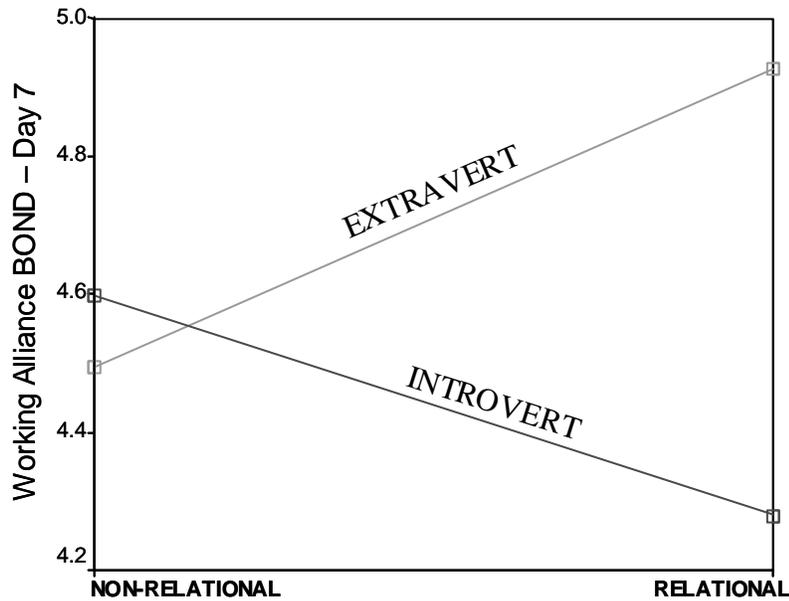


Figure 9-24. Effects of Relational Strategies on Introverts and Extraverts (not significant)

9.6.5 Generality

Finally, the generality of the findings deserves consideration. Although there were subjects in the study representing a fairly diverse population, the majority were students (69%) and well educated (59% had college degrees). However, comparisons between the students and non-students failed to yield any consistent pattern in either ability to change (measured by attitudes or behavior) or desire to continue working with FitTrack, and none of the differences observed were statistically significant. This provides some evidence that the FitTrack program could be successfully deployed for a much larger and wider audience that has access to home computers.

9.6.6 Laura version 2.0

As with the study in chapter 5, there are a number of important lessons learned in this effort that should be incorporated into future studies of this type. First and foremost, there were a number of problems in the software and experimental protocol that need to be corrected. Most important is the screening of subjects with respect to their baseline levels of activity. The initial 7-day recall should have been done as part of the intake interview (rather than on first login to the system), so that the experimenter could carefully explain how the form is filled out, define all terms used (e.g., “moderate activity”) and verify that subjects were under the 30 minute level before admitting them to the study.

As mentioned above, probably the most important things to change in a follow-on study are increases in the number of subjects and the length of the study. The complexity of the system (e.g., size of the dialogue network) could be maintained by simply requiring that subjects only access the system once or twice a week rather than daily, but over a much longer period of time.

Simpler instruments than the Working Alliance Inventory could be used to measure the quality of relationship over time (e.g., (Miell, 1987)), enabling them to be used more frequently to provide finer-grained information on the course of relationship development. This would enable investigations into the relational effects of specific dialogue exchanges, something not possible in the current study since the working alliance was only evaluated at two time points.

There are two control conditions that would be important to look at in the future. The first would be to compare the current experimental conditions with interactions with a human trainer, either using face-to-face meetings or using a protocol and technology similar to the one outlined in section 7.4.2. Another important comparison would be to investigate the effects of removing the animation from the current RELATIONAL condition, to determine what effects, if any, the presence and nonverbal behavior of the character's image have on subjects.

There are many other interesting questions to investigate within the FitTrack framework. For example, in the current system Laura never berates subjects for failing to perform. It would be interesting to include the use of negative feedback by the agent (or at least signs of frustration or disappointment) when users repeatedly fail to meet their goals, and see how this impacts performance. It would also be interesting to provide subjects with a much wider range of dialogue choices and track how much they talk to the agent and the specific topics they choose, and correlate this behavior with relational and outcome measures. Another question that arose during the study concerns the efficacy of computer-based empathy—which is most important for feeling “heard”: allowing a user to freely express themselves (e.g., via speech or free text input) or providing feedback that is tailored to their emotional state? Given the difficulty in inferring emotional state from natural language input (Liu, Lieberman, & Selker, 2003), this remains an important issue for the construction of systems that demonstrate artificial caring. Finally, investigations into ways of alleviating dialogue repetitiveness and determining its impact on attitudes would constitute important follow on work.

9.7 Summary

This chapter presented an evaluation of the use of relational face-to-face dialogue by an exercise advisor agent over repeated interactions with subjects, and its effect on subjects' reported relationship with the agent and their exercise behavior. The use of relational dialogue and nonverbal behavior did result in significant increases in a number of measures of the quality of the relationship, with subjects rating the agent significantly higher on likeability, trust, respect, feelings it cared for them, and willingness to continue interacting with it, relative to an agent that did not use these techniques. Although subjects in all experimental groups showed significant gains in exercise self-efficacy and behavior during the intervention part of the study, the effects of relationship building did not translate into exercise gains. Possible reasons for this include there being too few subjects in the study, the duration of the study being too short, and the need for additional interventions to work in conjunction with the relational techniques implemented, for example, to improve the agent's ability to negotiate the goals and tasks of the intervention with subjects.

Chapter 10

CONCLUSION

10.1 Summary of Significant Contributions

The preceding chapters have presented the motivation for the development of relational agents, and a series of explorations in the design, development, and evaluation of these agents.

The theory presented in Chapter 3 points the way to developing theoretical models of relationship based on theories of collaborative behavior in philosophy and Artificial Intelligence, and serves as a bridge between those fields and the social psychology of personal relationships. Although relational agents can be developed without explicit representation of the beliefs, desires and intentions of the agent and user, this level of representation may ultimately be required for agents that have the broadest degree of generality and adaptability.

The new model of face threat presented in Chapter 3 is novel within the fields of sociolinguistics and computational linguistics, in that it takes discourse context into account (i.e., the sequence of preceding utterances), rather than being based solely on isolated speech acts. The dialogue planner presented in Chapter 4 showed how this model could be implemented in a dialogue planner to produce relationally-appropriate sequencing of dialogue moves.

The software architecture presented in Chapter 6 demonstrates a novel, extensible framework for generating natural nonverbal behavior for an embodied relational agent based on the surface form of the utterance that the agent is to produce.

Chapters 7 and 8 presented several strategies that can be used by a relational agent to establish and maintain a relationship with a person over time, and some design guidelines for how these strategies can be implemented.

Chapter 8 presented an embodied conversational agent that incorporated many firsts. It is the first designed for long-term interactions with users, and which incorporates the ability to remember things about users between interactions and refer back to them in subsequent dialogues. It is the first embodied conversational agent designed for use on home computers that incorporates a wide range of naturalistic coverbal behavior, including hand gestures, posture shifts, and facial animation. It is the first embodied conversational agent designed for scalable client-server deployment to support a large number of users. It also includes the widest range of verbal and nonverbal behaviors yet developed for relationship-building and emotional support.

The two studies presented in Chapters 5 and 9 demonstrated that people will readily engage in relational dialogue with a software agent, and the second study demonstrated the positive impacts this had on subject's perceived relationship with the agent.

Finally, perhaps the most significant contribution of this thesis is an existence proof that relational agents can build trusting, caring relationships with people that can be sustained over multiple interactions.

10.2 Practical Implications

There are several lessons learned from this work that are applicable to the practical matters of building computer interfaces and conducting human subjects experiments.

10.2.1 Implications for Human-Computer Interface Design

One of the most intriguing findings from the FitTrack experiment is that, without any training or instruction, all 100 subjects managed to use the ECA interface without any problems, or without having their expectations dashed (a common criticism of anthropomorphic interfaces). Part of this is likely due to some principles of the social psychology of personal relationships that could be productively used in other kinds of software interfaces as well:

- Meta-relational communication – being very clear up front about the roles and expectations of each of the parties in the relationship, and checking in from time-to-time to see how everything is going and making adjustments as needed.
- Appropriate use of politeness and facework – this not only includes appropriate language forms when making potentially threatening requests of the user (instead of “Enter your social security number”, how about “Could you please enter your social security number”?) but avoiding violations of the user’s expectations in general (e.g., gradually transitioning into new tasks or topics).
- And, as noted by Klein (Klein, Moon, & Picard, 2002), appropriate use of empathy by a computer can go a long way towards making them feel understood and alleviating negative emotional states such as frustration.

Another important lesson learned for HCI is that, while reliability and consistency are highly prized in most aspects of interface design, there are some applications areas in which variability, even randomness, is important for keeping the user engaged in the task.

10.2.2 Implications for Experimental Design

The experimental protocol used in the FitTrack experiment opens up a whole new realm of longitudinal designs, in which experimental manipulations and their effects can be separated by days or weeks, and in which subjects’ perceptions of an interface agent can be evaluated the second, third, or twentieth time they use it—something that has been sadly missing in most evaluations of interface agents.

Dahlback wrote that wizard-of-oz studies are useful for collecting empirical data on the use of systems that had not been fully developed yet (Dahlback, Jonsson, & Ahrenberg, 1993). The study described in Chapter 5 presents another use for this kind of experimental design; the evaluation of theories of how face-to-face conversation works. In this experimental design: two or more alternative models of how conversation works (including verbal and/or nonverbal behavior) are used to drive the behavior of an ECA, with a human wizard factoring out all of the speech and natural language understanding problems from the equation, and the subject’s resultant behaviors and attitudes recorded and analyzed to determine which alternative best met the conversational goals of the agent.

10.3 Sociological Implications

As a technological optimist, I like to focus on the positive ways in which technologies can help society. The dream is that relational agents can help people to lead happier, healthier lives, not only by helping them eliminate harmful behaviors and adopt ones that lead to greater health and enjoyment, but also by putting a small dent in the 60% of the \$1.2 trillion per year on spent on healthcare in the U.S. that is attributable to behavioral problems (Prochaska, 2001).

On the dark side, relational agents could conceivably lead to further fractionating of society if, rather than supplying additional social bonds they tend to replace the ones that people already have, or would have had, with other people. However, relational agents could play a positive role in socialization. One way is by acting as social role models. In developing FitTrack we joked that it could actually teach socially-backward MIT students to conduct social dialogue, and in fact one subject in the Rea study even mentioned this as a possibility after using the system (“I was thinking that if she can do it, then any person can learn how to chit chat.”). Another way that relational agents could actually increase socialization is by providing social network support. Imagine if, after a series of set backs at work, your agent contacts your best friend on your behalf, tells them what is going on and arranges an outing for you. Alternatively, your agent could introduce you to a support group of people who are currently going through similar problems.

Another sociological concern is that if we come to rely too much on our agents for our relational needs, this may end up “dumbing us down” to their level (Lanier), potentially producing a race of social simpletons. The answer here is to both ensure that these agent interfaces are as rich as possible, and ensure that we use them to augment human relationships rather than replace them.

10.4 Research Implications

10.4.1 Psychology

One of the great potentials of this technology that has not been discussed yet is its use as a tool for understanding human relationships. In the hands of social psychologists, agents like Laura could be used to very subtly and precisely manipulate different aspects of user-agent interactions over time (e.g., testing a particular maintenance behavior or the withholding of an empathetic response) to see what the long-term effects on their relationship are.

10.4.2 Linguistics

The pervasiveness of phenomena such as social deixis, politeness, common ground and “phaticity” in everyday relational language serves as an additional reminder that the syntax-centric orientation of contemporary linguistics needs to be expanded to include a wider range of pragmatic phenomena. On the plus side, this thesis has shown that there may indeed be application areas in which the proper strategic and tactical uses of language by a computer can make a positive difference in task outcomes, not to mention that relational language (embodied and otherwise) offers up a plethora of new research problems for linguists to sink their teeth into.

10.4.3 Artificial Intelligence

There has been a fair amount of work over the last few years on models of negotiation, cooperation and even “relationship” in the sub-field of Distributed Artificial Intelligence, exemplified by the work of Marsh on models of inter-agent trust (discussed in Chapter 3)(Marsh, 1994). However, little of this work has been applied to understanding and modeling the relationships between people and computer agents; thus, this represents an interesting area of future research. In addition, as suggested by the dialogue planning system presented in Chapter 4, planning human relational behavior can provide the field of AI with another challenging domain to work in, given all the complexities of trading off multiple social, task and linguistic goals. A final lesson learned here is that complex real-time reasoning and dialogue planning are not always needed, even for sustained interactions with a seemingly intelligent agent. In the FitTrack study, an essentially scripted interaction, devoid of any complex inferencing or planning capability, was sufficient to help people through a month of interactions. Of course, if the agent were to be re-implemented for another domain it would require an extensive amount of research and scripting.

10.5 Ethical Considerations

Relational agents, as any technology, can be abused. Agents which earn our trust over time can be used to provide more potent means of persuasion for marketers than more passive forms of advertising. If we eventually come to rely on our agents as sources of grounding for our beliefs, values and emotions (one of the major functions of close human relationships (Duck, 1991)) then they could become a significant source of manipulation and control over individuals or even over entire societies.

There are those who also feel that any anthropomorphic interface is unethical, because it unrealistically raises users’ expectations. One way to combat this problem is through proper meta-relational communication—having the agent be as clear as possible about what it can and can’t do, and what expectations the user should have about their respective roles in the interaction. For example, in her first interaction with users, Laura states:

I’m going to help you meet your exercise goals over the next month. I’m going to be your exercise advisor. My job is to help you set goals and overcome obstacles, and I’ll also be giving you lots of tips on how to exercise. You need to keep in mind that I’m just a computer character with limited capabilities, so I hope you will understand if there are some things I cannot help you with.

Such language can help users properly adjust their expectations at the start of an interaction.

I also believe that avoiding free form natural language input, whether it be speech or text, is possibly the single best way to avoid unrealistically raising expectations in anthropomorphized interfaces. Dynamic menus were chosen for the FitTrack application because of their reliability and simplicity, but in retrospect they also did an excellent job of communicating to the user exactly which responses the agent was capable of handling at any point in the conversation. Until the AI and NLP problems have been solved, offering the user the option of saying or typing anything in an unconstrained context, when the system can actually only handle a handful of responses, is at least bad interface design and at most dishonest.

A final issue, raised by Picard and Klein, is the ethic of building agents that pretend to care, understand and empathize, when, in fact, they have no emotions of their own (advanced

systems with the analogue of neurological and emotional systems are another story) (Picard & Klein, 2002). As observed by Turkle, people today seem quite comfortable with computational artifacts that only appear to have emotions (Turkle, 1995) and, as confirmed by most users in my FitTrack study, the end seems to justify the means. As one subject put it:

She's a computer character. I don't know if she cared about me. I don't know if she feels. She's a character and has a role, but I don't know if she has feelings. But, it worked for me and I'm happy.

10.6 Future Directions

There are many, many fruitful directions this research program could be advanced in the future. There are interesting research problems presented in each of the chapters in this thesis:

- *Chapter 3* – A full theory of relationship based on belief-desire-intention models of collaborative behavior remains a challenge to develop. In particular, the development of a calculus for relational scorekeeping and collaborative decision making may be required before relational agents could make use of these models.
- *Chapter 4* – The discourse planner for mixed task and social dialogue awaits integration with a dialogue planning framework that brings in more of the intentional and attentional structures of discourse context (like COLLAGEN (Rich, Sidner, & Lesh, 2001) or DPOCL (Young & Moore, 1994)) to increase the capability and generality of the system. The planner could also benefit from integration with a true natural language text generation system to increase its flexibility and provide a partial solution to the problem of relational agent repetitiveness discussed in Chapter 9.
- *Chapter 5* – An entire series of wizard-of-oz experiments simulating face-to-face conversation with a relational agent could be performed in which various aspects of the relational model are manipulated—such as nonverbal behavior, presumed relevance of topics, topic coherence, face threat parameters, etc.—rather than just the inclusion or non-inclusion of small talk. Additional measures of trust and face threat that are more sensitive to immediate context should be used, such as subject gaze behavior or galvanic skin response. If task outcomes are of interest, then a task domain other than real estate sales should be explored.
- *Chapter 6* – There are innumerable studies on the role of nonverbal behavior in face-to-face conversation that could be done to further extend the BEAT text-to-embodied-speech system. Of particular relevance to relational agents are studies that can be used to derive quantitative models of nonverbal behavior signifying different conversational frames and relational stances. Equally important is work on sensing and pattern recognition technologies that can sense and interpret user nonverbal behavior in relational interactions.
- *Chapter 7* – The relational agent described in Chapters 8 and 9 represents a single point in a large space of helping applications. There are potentially many other helping domains that could benefit from the deployment of a relational agent, from other health behavior change domains (e.g., smoking cessation, diet) to coaching, counseling and therapy. Even within health behavior change, there may be other “processes of change” (Prochaska, Norcross, & Diclemente, 1994) (e.g., consciousness raising, stimulus control, and reinforcement) which could be facilitated

by a relational agent, in addition to providing a “helping relationship”. Examples of other behavior change strategies that could be deployed include training for relapse prevention and other maintenance techniques that help subjects continue a desired behavior change after the intervention has ended.

- *Chapter 8* – Although the scripting language for the FitTrack relational agent was very sophisticated, ultimately it should be replaced with a natural language text generation system (as described above). Interesting research problems for long-term relationships relative to this change include how to refer back to prior conversations (what does a historical discourse context look like?) and how to incrementally populate such a system with new knowledge and topics of conversation so that someone could use such a system for an indefinite period of time without it repeating itself. Such a system should also be capable of determining when its relational strategies are not working (perhaps over many interactions) and changing these strategies accordingly.
- *Chapter 9* – The FitTrack study should be repeated over a longer period of time with a much larger population, and one that is more representative of typical sedentary adults in the U.S. The dialogue strategies should also be extended to include detailed negotiations of the tasks and goals of the intervention, and the interactions should be designed to gradually “wean” subjects off the agent at the end of the intervention. One interesting research question that arose from the study is exactly how much change and what kinds of change are needed in an agent’s language for users to avoid perceiving it as repetitive? Another interesting question concerns determining the input and output modalities required for effective empathic understanding. Simpler measures than the Working Alliance Inventory could be used to measure the quality of relationship over time (e.g., (Miell, 1987)), enabling them to be used more frequently to provide finer-grained information on the course of relationship development. And, of course, the FitTrack architecture and experimental protocol could be used for other health behavior change studies, such as smoking cessation.

Finally, relational agents on mobile devices could provide a particularly powerful combination, both for relationship-building (a “buddy” who is always with you) and for behavior change (e.g., providing interventions at the time and place of need, having a workout coach that you can take to the gym, etc.). The initial conception for FitTrack was a health advisor on a PDA, motivating a study of how people would interact with ECAs on handheld devices (see Figure 9-1) (Bickmore, 2002). This remains a fruitful area for further research.



Figure 10-1. Handheld ECA

10.7 Concluding Remarks

This thesis defined relational agents, demonstrated that they can be developed and deployed on home computers, demonstrated that they are capable of significant bonding with users in ways that significantly affect their attitudes and behavior, and further demonstrated uses of these agents in sales and health behavior change application domains.

Given this, this work represents one of the first steps in the development of technologies for providing social support to people who may not otherwise have the benefit of a caring network of family, friends, and professional helpers, or augmenting the social support networks of those who do. For applications in which the primary objective is not social support, this work provides a partial roadmap of social competencies required of computer agents designed to interact with users in social contexts or to have long-term roles in their lives. For science, this work provides new tools and paradigms for investigating and thinking about human relationships.

This technology has the potential to have a negative effect on society by further eroding existing social bonds between people, but with the proper stewardship it has great potential for transforming and liberating society as well.

References

- Allison, K., Dwyer, J., & Makin, S. (1999). Perceived Barriers to Physical Activity among High School Students. *Preventive Medicine*, 28, 608-615.
- Altman, I., & Taylor, D. (1973). *Social penetration: The development of interpersonal relationships*. New York: Holt, Rinhart & Winston.
- Andersen, P., & Guerrero, L. (1998). Principles of Communication and Emotion in Social Interaction. In P. Andersen & L. Guerrero (Eds.), *Handbook of Communication and Emotion*. New York: Academic Press.
- Andersen, P., & Guerrero, L. (1998a). The Bright Side of Relational Communication: Interpersonal Warmth as a Social Emotion. In P. Andersen & L. Guerrero (Eds.), *Handbook of Communication and Emotion* (pp. 303-329). New York: Academic Press.
- Andersen, P., & Guerrero, L. (1998b). Principles of Communication and Emotion in Social Interaction. In P. Andersen & L. Guerrero (Eds.), *Handbook of Communication and Emotion*. New York: Academic Press.
- Andre, E., Muller, J., & Rist, T. (1996). The PPP Persona: A Multipurpose Animated Presentation Agent. Paper presented at the Advanced Visual Interfaces.
- Andre, E., Rist, T., & Muller, J. (1998). Integrating reactive and scripted behaviors in a life-like presentation agent. *Proceedings of AGENTS'98*, 261-268.
- Anselmi, K., & James E. Zemanek, J. (1997). Relationship Selling: How Personal Characteristics of Salespeople Affect Buyer Satisfaction. *Journal of Social Behavior and Personality*, 12(2), 539-550.
- Argyle, M. (1988). *Bodily Communication*. New York: Methuen & Co. Ltd.
- Bachelor, A. (1991). Comparison and Relationship to Outcome of Diverse Dimensions of the Helping Alliance as Seen by Client and Therapist. *Psychotherapy*, 28(4), 534-549.
- Bagchi, S., Biswas, G., & Kawamura, K. (1996). Interactive task planning under uncertainty and goal changes. *Robotics and Autonomous Systems*, 18.
- Ball, G., & Breese, J. (2000). Emotion and Personality in a Conversational Agent. In J. Cassell & J. Sullivan & S. Prevost & E. Churchill (Eds.), *Embodied Conversational Agents*. Cambridge, MA: MIT Press.
- Barber, B. (1983). *Logic and Limits of Trust*. New Jersey: Rutgers University Press.
- Bassett, D., Cureton, A., & Ainsworth, B. (2000). Measurement of daily walking distance-questionnaire versus pedometer. *Medicine & Science in Sports & Exercise*, 32(5), 1018-1023.
- Bateson, G. (1954). *A theory of play and fantasy. Steps to an ecology of mind*. New York: Ballantine.

- Berscheid, E., & Reis, H. (1998). Attraction and Close Relationships. In D. Gilbert & S. Fiske & G. Lindzey (Eds.), *The Handbook of Social Psychology* (pp. 193-281). New York: McGraw-Hill.
- Beskow, J., & McGlashan, S. (1997). Olga: a conversational agent with gestures. Paper presented at the IJCAI 97.
- Beutler, L. E., Crago, M., & Arizmendi, T. G. (1986). Therapist Variables in Psychotherapy Process and Outcome. In A. E. Bergin & S. L. Garfield (Eds.), *Handbook of Psychotherapy and Behavior Change* (pp. 257-310). New York: John Wiley & Sons.
- Bickmore, T. (2002). Towards the Design of Multimodal Interfaces for Handheld Conversational Characters. Paper presented at the Human Factors in Computing Systems (CHI), Minneapolis, MN.
- Bickmore, T., & Cassell, J. (2001). Relational Agents: A Model and Implementation of Building User Trust. Paper presented at the Proceedings of CHI'01, Seattle, WA.
- Bickmore, T., & Cassell, J. (to appear). Social Dialogue with Embodied Conversational Agents. In J. van Kuppevelt & L. Dybkjaer & N. Bernsen (Eds.), *Natural, Intelligent and Effective Interaction with Multimodal Dialogue Systems*. New York: Kluwer Academic.
- Binstead, K. (1998). Designing portable characters. Paper presented at the Workshop on Embodied Conversational Characters, Lake Tahoe, CA.
- Bowlby, J. (1979). *The making and breaking of affectional bonds*. London: Tavistock.
- Bratman, M. (1990). What is Intention? In P. Cohen & J. Morgan & M. Pollack (Eds.), *Intentions in Communication* (pp. 15-31). Cambridge, MA: MIT Press.
- Breazeal, C. (2000). *Sociable Machines: Expressive Social Exchange Between Humans and Robots*. Unpublished Ph.D., MIT, Cambridge.
- Brehm, S. (1992). *Intimate Relationships*. New York: McGraw-Hill.
- Brown, P., & Levinson, S. (1978). Universals in language usage: Politeness phenomena. In E. Goody (Ed.), *Questions and Politeness: Strategies in Social Interaction*. Cambridge: Cambridge University Press.
- Brown, P., & Levinson, S. C. (1987). *Politeness: Some universals in language usage*. Cambridge: Cambridge University Press.
- Brown, R., & Gilman, A. (1972). The pronouns of power and solidarity. In P. Giglioli (Ed.), *Language and Social Context* (pp. 252-282.). Harmondsworth: Penguin.
- Brown, R., & Gilman, A. (1989). Politeness theory and Shakespeare's four major tragedies. *Language in Society*, 18, 159-212.
- Buck, R. (1991). Social Factors in Facial Display and Communication: a Reply to Chovil and Others. *Journal of Nonverbal Behavior*, 15(3), 155-167.
- Buller, D., & Street, R. (1992). Physician-Patient Relationships. In R. Feldman (Ed.), *Application of Nonverbal Behavioral Theories and Research* (pp. 119-141). Hillsdale, NJ: Lawrence Erlbaum.

- Burgoon, J. K., & Hale, J. L. (1984). The Fundamental Topoi of Relational Communication. *Communication Monographs*, 51, 193-214.
- Burleson, B., & Goldsmith, D. (1998). How the Comforting Process Works: Alleviating Emotional Distress through Conversationally Induced Reappraisals. In P. Andersen & L. Guerrero (Eds.), *Handbook of Communication and Emotion*. New York: Academic Press.
- Cahn, J. (1990). The Generation of Affect in Synthesized Speech. *Journal of the American Voice I/O Society*, 8, 1-19.
- Cassell, J. (1999). Embodied Conversation: Integrating Face and Gesture into Automatic Spoken Dialogue systems. In S. Luperfoy (Ed.), *Spoken Dialogue Systems*. Cambridge: MIT Press.
- Cassell, J., & Bickmore, T. (2000). External Manifestations of Trustworthiness in the Interface. *Communications of the ACM*, 43(12), 50-56.
- Cassell, J., & Bickmore, T. (2001). Negotiated Collusion: Modeling Social Language and its Relationship Effects in Intelligent Agents. *User Modeling and Adaptive Interfaces*, in press.
- Cassell, J., & Prevost, S. (1996). Distribution of Semantic Features Across Speech and Gesture by Humans and Computers. *Proceedings of the Workshop on the Integration of Gesture in Language and Speech*.
- Cassell, J., Ananny, M., Basu, A., Bickmore, T., Chong, P., Mellis, D., Ryokai, K., Smith, J., Vilhjálmsón, H., & Yan, H. (2000). Shared Reality: Physical Collaboration with a Virtual Peer. Paper presented at the Proceedings of CHI '00, The Hague, The Netherlands.
- Cassell, J., Bickmore, T., Billinghamurst, M., Campbell, L., Chang, K., Vilhjálmsón, H., & Yan, H. (1999). Embodiment in Conversational Interfaces: Rea. Paper presented at the Proceedings of CHI '99, Pittsburgh, PA.
- Cassell, J., Bickmore, T., Vilhjálmsón, H., & Yan, H. (2000). More Than Just a Pretty Face: Affordances of Embodiment. Paper presented at the IUI 2000, New Orleans, Louisiana.
- Cassell, J., Nakano, Y., Bickmore, T., Sidner, C., & Rich, C. (2001). Non-Verbal Cues for Discourse Structure. Paper presented at the Association for Computational Linguistics.
- Cassell, J., Pelachaud, C., Badler, N., Steedman, M., Achorn, B., Becket, T., Douville, B., Prevost, S., & Stone, M. (1994). Animated Conversation: Rule-based Generation of Facial Expression, Gesture and Spoken Intonation for Multiple Conversational Agents. Paper presented at the SIGGRAPH, Orlando.
- Cassell, J., Stocky, T., Bickmore, T., Gao, Y., Nakano, Y., Ryokai, K., Tversky, D., Vaucelle, C., & Vilhjálmsón, H. (2002, February 12-15). MACK: Media lab Autonomous Conversational Kiosk. Paper presented at the Imagina '02, Monte Carlo.
- Cassell, J., Sullivan, J., Prevost, S., & Churchill, E. (2000). *Embodied Conversational Agents*. Cambridge: MIT Press.
- Cassell, J., Vilhjálmsón, H., & Bickmore, T. (2001). BEAT: The Behavior Expression Animation Toolkit. Paper presented at the SIGGRAPH '01, Los Angeles, CA.

- Cegala, D., Waldro, V., Ludlum, J., McCabe, B., Yost, S., & Teboul, B. (1988). A study of interactants' thoughts and feelings during conversation. Paper presented at the Ninth Annual Conference on Discourse Analysis, Philadelphia, PA.
- Celio, A., Winzelberg, A., Dev, P., & Taylor, C. (2002). Improving Compliance in On-line Structured Self-help Programs: Evaluation of an Eating Disorder Program. *Journal of Psychiatric Practice*, 8(1), 14-20.
- Cheepen, C. (1988). *The Predictability of Informal Conversation*. New York: Pinter.
- Chisolm, D., Collis, M., Kulak, L., Davenport, W., & Gruber, N. (1975). Physical activity readiness. *British Columbia Medical Journal*, 17, 375-378.
- Clark, H. H. (1992). *Arenas of Language Use*. Chicago, IL: University of Chicago Press.
- Clark, H. H. (1996). *Using Language*. Cambridge: Cambridge University Press.
- Colby, K. (1981). Modeling a paranoid mind. *The Behavioral and Brain Sciences*, 4(4), 515-560.
- Cole, T., & Bradac, J. (1996). A Lay Theory of Relational Satisfaction with Best Friends. *Journal of Social and Personal Relationships*, 13(1), 57-83.
- Cole, T., & Bradac, J. (1996). A Lay Theory of Relational Satisfaction with Best Friends. *Journal of Social and Personal Relationships*, 13(1), 57-83.
- Connors, G., Carroll, K., DiClemente, C., & Longabaugh, R. (1997). The Therapeutic Alliance and Its Relationship to Alcoholism Treatment Participation and Outcome. *Journal of Consulting and Clinical Psychology*, 65(4), 588-598.
- Coupland, J., Coupland, N., & Robinson, J. D. (1992). "How are you?": Negotiating phatic communion. *Language in Society*, 21, 207-230.
- Csikszentmihalyi, M., & Rochberg-Halton, E. (1998). *The meaning of things: Domestic symbols and the self*. Cambridge: Cambridge University Press.
- Dahlback, N., Jonsson, A., & Ahrenberg, L. (1993). Wizard of Oz Studies: Why and How. Paper presented at the IUI 93.
- Dahlback, N., Jonsson, A., & Ahrenberg, L. (1998). Wizard of Oz Studies -- Why and How. In M. T. Maybury & W. Wahlster (Eds.), *Readings in Intelligent User Interfaces* (pp. 610-619). San Francisco, CA: Morgan Kaufmann Publishers, Inc.
- Dainton, M., & Stafford, L. (1993). Routine Maintenance Behaviors: A comparison of relationships type, partner similarity and sex differences. *Journal of Social and Personal Relationships*, 10, 255-271.
- Damon, W., & Phelps, E. (1989). Strategic Uses of Peer Learning in Children's Education. In T. Berndt & G. Ladd (Eds.), *Peer Relationships in Child Development* (pp. 135-157). New York: Wiley.
- Dehn, D. M., & Mulken, S. v. (2000). The Impact of Animated Interface Agents: A Review of Empirical Research. *International Journal of Human-Computer Studies*, 52, 1-22.

- Deutsch, M. (1973). *The Resolution of Conflict: Constructive and Destructive Processes*. New Haven: Yale University Press.
- Devitt, M., & Sterelny, K. (1999). *Language and Reality: An Introduction to the Philosophy of Language*. Cambridge, MA: MIT Press.
- Dindia, K. (1994). A Multiphasic View of Relationship Maintenance Strategies. In D. Canary & L. Stafford (Eds.), *Communication and Relational Maintenance* (pp. 91-112). New York: Academic Press.
- Dishman, R. (1994). Introduction: Consensus, Problems and Prospects. In R. Dishman (Ed.), *Advances in Exercise Adherence* (pp. 1-27). Champaign, IL: Human Kinetics.
- Dishman, R., & Sallis, J. (1994). Determinants and Interventions for Physical Activity and Exercise. In C. Bouchard & R. Shephard & T. Stephens (Eds.), *Physical Activity, Fitness and Health* (pp. 214-238). Champaign, IL: Human Kinetics.
- Drexler, M. (1999). Matsushita releases electronic pet. IDG.net. March 24. <http://www.idg.net/idgns/1999/03/24/MatsushitaReleasesElectronicPet.shtml>
- Drummond, M., & Bresina, J. (1990). Anytime synthetic projection: Maximizing the probability of goal satisfaction, *AAAI-90* pp. 138-144.
- Duck, S. (1988). *Relating to others*. Chicago: Dorsey.
- Duck, S. (1991). *Understanding Relationships*. New York: Guilford Press.
- Duck, S. (1994). *Meaningful Relationships*. London: SAGE Publications.
- Duck, S. (1995). Talking Relationships into Being. *Journal of Social and Personal Relationships*, 12(4), 535-540.
- Duck, S. (1998). *Human Relationships*. London: SAGE Publications.
- Dunbar, R. (1996). *Grooming, Gossip, and the Evolution of Language*. Cambridge, MA: Harvard University Press.
- Duncan, S. (1974). On the structure of speaker-auditor interaction during speaking turns. *Language in Society*, 3, 161-180.
- Ekman, P. (1993). Facial expression and emotion. *American Psychologist*, 48(4), 384-392.
- Elliott, C. (1992). *The Affective Reasoner: A process model of emotions in a multi-agent system*. Northwestern University.
- Fikes, R., & Nilsson, N. (1971). STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial Intelligence*, 5(2), 189-208.
- Fillmore, C. (1975). Pragmatics and the description of discourse. In P. Cole (Ed.), *Radical pragmatics* (pp. 143-166). New York: Academic Press.
- Fillmore, C. (1976). The need for a frame semantics within linguistics, *Statistical methods in in linguistics* (pp. 5-29). Stockholm: Skriptor.

- Finin, T., Weber, J., Wiederhold, G., Genesereth, M., Fritzson, R., McKay, D., McGuire, J., Pelavin, R., Shapiro, S., & Beck, C. (1994). Specification of the KQML Agent-Communication Language.
- Frankel, R. (1995). Emotion and the Physician-Patient Relationship. *Motivation and Emotion*, 19(3), 163-173.
- Freud, S. (1913). On the beginning of treatment: Further recommendations on the technique of psychoanalysis. In J. Strachey (Ed.), *Standard edition of the complete psychological works of Sigmund Freud* (pp. 122-144). London: Hogarth Press.
- Fukayama, A., Ohno, T., Mukawa, N., Sawaki, M., & Hagita, N. (2002). Messages Embedded in Gaze of Interface Agents: Impression management with agent's gaze. Paper presented at the CHI'02, Minneapolis, MN.
- Gabarro, J. (1990). The Development of Working Relationships. In J. Galegher & R. Kraut & C. Egidio (Eds.), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work* (pp. 79-110). Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- Gaston, L. (1990). The Concept of the Alliance and its Role in Psychotherapy: Theoretical and Empirical Considerations. *Psychotherapy*, 27(2), 143-153.
- Gelso, C., & Hayes, J. (1998). *The Psychotherapy Relationship: Theory, Research and Practice*. New York: John Wiley and Sons.
- Gilbertson, J., Dindia, K., & Allen, M. (1998). Relational Continuity Constructional Units and the Maintenance of Relationships. *Journal of Social and Personal Relationships*, 15(6), 774-790.
- Gill, D., Christensen, A., & Fincham, F. (1999). Predicting marital satisfaction from behavior: Do all roads really lead to Rome? *Personal Relationships*, 6, 369-387.
- Goetz, P. (1997). *Attractors in Recurrent Behavior Networks*. Unpublished PhD, State University of New York at Buffalo, Buffalo, NY.
- Goffman, E. (1983). *Forms of Talk*. Philadelphia, PA: University of Pennsylvania Publications.
- Goffman, I. (1967). On face-work, *Interaction Ritual: Essays on Face-to-Face Behavior* (pp. 5-46). New York: Pantheon.
- Goldsmith, D., McDermoot, V., & Alexander, S. (2000). Helpful, supportive and sensitive: Measuring the evaluation of enacted social support in personal relationships. *Journal of Social and Personal Relationships*, 17(3), 369-391.
- Goleman, D. (1995). *Emotional Intelligence*. New York: Bantam Books.
- Goodwin, C. (1981). Achieving Mutual Orientation at Turn Beginning, *Conversational Organization: Interaction between Speakers and Hearers* (pp. 55-89). New York: Academic Press.
- Goodwin, M. (1996). Shifting Frame. In D. Slobin & J. Gerhardt & A. Kyratzis & J. Guo (Eds.), *Social Interaction, Social Context, and Language: Essays in Honor of Susan Ervin-Tripp*.

- Grice, P. (1989). *Studies in the Way of Words*. Cambridge, MA: Harvard University Press.
- Grosz, B. J., & Kraus, S. (1996). Collaborative Plans for Complex Group Action. *Artificial Intelligence*, 86(2), 269-357.
- Grosz, B., & Kraus, S. (1996). Collaborative plans for complex group action. *Artificial Intelligence*, 86(2), 269--357.
- Grosz, B., & Sidner, C. (1986). Attention, Intentions, and the Structure of Discourse. *Computational Linguistics*, 12(3), 175-204.
- Grosz, B., & Sidner, C. (1990). Plans for Discourse. In P. R. Cohen & J. Morgan & M. E. Pollack (Eds.), *Intentions in Communication* (pp. 417-444). Cambridge, MA: MIT Press.
- Gumperz, J. (1977). Sociocultural Knowledge in Conversational Inference. In M. Saville-Troike (Ed.), *Linguistics and Anthropology* (pp. 191-211). Washington DC: Georgetown University Press.
- Gutek, B., Cherry, B., Bhappu, A., Schneider, S., & Woolf, L. (2000). Features of Service Relationships and Encounters. *Work and Occupations*, 27(3), 319-352.
- Halliday, M. (1973). *Explorations in the Functions of Language*. London: Edward Arnold.
- Hanks, S. (1994). *Discourse Planning: Technical Challenges for the Planning Community*. Paper presented at the AAAI Workshop on Planning for Inter-Agent Communication.
- Hartup, W. (1996). Cooperation, close relationships, and cognitive development. In W. Bukowski & A. Newcomb & W. Hartup (Eds.), *The company they keep: Friendship in childhood and adolescence* (pp. 213-237). Cambridge: Cambridge University Press.
- Havens, L. (1986). *Making Contact: Uses of Language in Psychotherapy*. Cambridge, MA: Harvard University Press.
- Heise, D. R. (1992). Affect Control Theory and Impression Formation. In M. Borgatta (Ed.), *Encyclopedia of Sociology* (Vol. I, pp. 12-17). New York: Macmillan.
- Henry, W., & Strupp, H. (1994). The Therapeutic Alliance as Interpersonal Process. In A. Horvath & L. Greenberg (Eds.), *The Working Alliance: Theory, Research and Practice* (pp. 51-84). New York: John Wiley & Sons.
- Hirschberg, J. (1990). Accent and Discourse Context: Assigning Pitch Accent in Synthetic Speech. Paper presented at the AAAI 90.
- Hiyakumoto, L., Prevost, S., & Cassell, J. (1997). Semantic and Discourse Information for Text-to-Speech Intonation. Paper presented at the ACL Workshop on Concept-to-Speech Generation, Madrid.
- Horvath, A. (1994). Research on the Alliance. In A. Horvath & L. Greenberg (Eds.), *The Working Alliance: Theory, Research and Practice* (pp. 259-286). New York: John Wiley & Sons.
- Horvath, A., & Greenberg, L. (1989). Development and Validation of the Working Alliance Inventory. *Journal of Counseling Psychology*, 36(2), 223-233.

- Horvath, A., & Luborsky, L. (1993). The Role of the Therapeutic Alliance in Psychotherapy. *Journal of Consulting and Clinical Psychology*, 61(4), 561-573.
- Horvath, A., & Symonds, B. (1991). Relation Between Working Alliance and Outcome in Psychotherapy: A Meta-Analysis. *Journal of Counseling Psychology*, 38(2), 139-149.
- Hovy, E. (1986). Putting Affect into Text. Paper presented at the Eighth Annual Conference of the Cognitive Science Society.
- Huang, X., Acero, A., Adcock, J., Hon, H.-W., Goldsmith, J., Liu, J., & Plumpe, M. (1996). Whistler: A Trainable Text-to-Speech System. Paper presented at the 4th Int'l Conf. on Spoken Language Processing (ICSLP '96), Piscataway, NJ.
- Izard, C. E. (1993). Four systems for emotion activation: Cognitive and noncognitive processes. *Psychological Review*, 100(1), 68-90.
- Jackendoff, R. (1999). The Natural Logic of Rights and Obligations. In R. Jackendoff & P. Bloom & K. Wynn (Eds.), *Language, Logic, and Concepts*. Cambridge, MA: MIT Press.
- Jacobson, R. (1960). Closing statement: linguistics and poetics. In T. A. Sebeok (ed.), *Style in Language*, Cambridge, MA: MIT Press., pp. 350-77.
- Jaworski, A., & Coupland, N. (1999). *The Discourse Reader*. London: Routledge.
- Jefferson, G. (1978). Sequential aspects of storytelling in conversation. In J. Schenkein (Ed.), *Studies in the organization of conversational interaction* (pp. 219-248). New York: Academic Press.
- Johnson-George, C., & Swap, W. (1982). Measurement of Specific Interpersonal Trust: Construction and Validation of a Scale to Assess Trust in a Specific Other. *J. Personality and Social Psychology*, 43(6), 1306-1317.
- Jonas, S. (1995). *Regular Exercise: A Handbook for Clinical Practice*. New York: Springer.
- Kapoor, A., & Picard, R. W. (2002). Real-Time, Fully Automatic Upper Facial Feature Tracking. Paper presented at the 5th International Conference on Automatic Face and Gesture Recognition, Washington, D.C.
- Kapoor, A., Mota, S., & Picard, R. W. (2001). Towards a Learning Companion that Recognizes Affect. Paper presented at the The American Association of Artificial Intelligence Conference, Falmouth, MA.
- Keijsers, G. P. J., Schaap, C. P. D. R., & Hoogduin, C. A. L. (2000). The Impact of Interpersonal Patient and Therapist Behavior on Outcome in Cognitive-Behavior Therapy. *Behavior Modification*, 24(2), 264-297.
- Kelley, H. (1983). Epilogue: An essential science. In H. Kelley & A. Berscheid & J. Christensen & T. Harvey & G. Huston & E. Levinger & L. McClintock & L. Peplau & D. Peterson (Eds.), *Close Relationships* (pp. 486-503). New York: Freeman.
- Kiesler, D. (1982). Interpersonal Theory for Personality and Psychotherapy. In J. Anchin & D. Kiesler (Eds.), *Handbook of Interpersonal Psychotherapy* (pp. 3-24). New York: Pergamon Press.

- Kiesler, S., & Sproull, L. (1997). 'Social' Human-Computer Interaction. In B. Friedman (Ed.), *Human Values and the Design of Computer Technology* (pp. 191-199). Stanford, CA: CSLI Publications.
- Kim, J., & Moon, J. Y. (1997). Emotional Usability of Customer Interfaces: Focusing on CyberBanking Systems Interfaces. Paper presented at the CHI 97.
- King, A. (2001). The Coming of Age of Behavioral Research in Physical Activity. *Annals of Behavioral Medicine*, 23(4), 227-228.
- Klein, J. T. (1999). *Computer Response to User Frustration* (Master's Thesis): MIT.
- Klein, J., Moon, Y., & Picard, R. (2002). This Computer Responds to User Frustration: Theory, Design, Results, and Implications. *Interacting with Computers*, 14, 119-140.
- Knapp, D. (1988). Behavioral Management Techniques and Exercise Promotion. In R. Dishman (Ed.), *Exercise Adherence: Its Impact on Public Health* (pp. 203-235). Champaign, Illinois: Human Kinetics Books.
- Koda, T., & Maes, P. (1996). Agents with Faces: The Effect of Personification. Paper presented at the Proceedings of the Fifth IEEE International Workshop on Robot and Human Communication (RO-MAN '96).
- Lanier, J. (1995). Agents of Alienation. *Interactions of the ACM*, 2(3), 66-72.
- Laver, J. (1975). Communicative functions of phatic communion. In A. Kendon & R. Harris & M. Key (Eds.), *The organization of behavior in face-to-face interaction* (pp. 215-238). The Hague: Mouton.
- Laver, J. (1981). Linguistic routines and politeness in greeting and parting. In F. Coulmas (Ed.), *Conversational routine* (pp. 289-304). The Hague: Mouton.
- Lazarus, R. (1994). Appraisal: The long and the short of it. In P. Ekman & R. J. Davidson (Eds.), *The nature of emotion: Fundamental questions* (pp. 208-215). New York: Oxford University Press.
- Lee, J., Kim, J., & Moon, J. (2000). What makes Internet users visit cyber stores again? Key design factors for customer loyalty. Paper presented at the CHI'00.
- Lester, J. C., Converse, S. A., Kahler, S. E., Barlow, S. T., Stone, B. A., & Bhogal, R. S. (1997). The Persona Effect: Affective Impact of Animated Pedagogical Agents. Paper presented at the Proceedings of CHI '97.
- Lester, J. C., Voerman, J. L., Towns, S. G., & Callaway, C. B. (1997). Cosmo: A Life-like Animated Pedagogical Agent with Deictic Believability. Paper presented at the IJCAI 97.
- Lester, J., Stone, B., & Stelling, G. (1999). Lifelike Pedagogical agents for Mixed-Initiative Problem Solving in Constructivist Learning Environments. *User Modeling and User-Adapted Interaction*, 9(1-2), 1-44.
- Lester, J., Towns, S., Callaway, C., Voerman, J., & Fitzgerald, P. (2000). Deictic and Emotive Communication in Animated Pedagogical Agents. In J. Cassell (Ed.), *Embodied Conversational Agents*. Cambridge, MA: MIT Press.

- Levinson, S. C. (1983). *Pragmatics*. Cambridge: Cambridge University Press.
- Lewis, R.A. (1972). A developmental framework for the analysis of premarital dyadic formation. *Family Process*, 11, 17-48.
- Lim, T. (1994). Facework and Interpersonal Relationships. In S. Ting-Toomey (Ed.), *The challenge of facework: Cross-cultural and interpersonal issues* (pp. 209-229). Albany, NY: State University of New York Press.
- Liu, H., Lieberman, H., & Selker, T. (2003). A Model of Textual Affect Sensing using Real-World Knowledge. Paper presented at the International Conference on Intelligent User Interfaces, Miami, Florida.
- Luborsky, L. (1994). Therapeutic Alliances as Predictors of Psychotherapy Outcomes: Factors Explaining the Predictive Success. In A. Horvath & L. Greenberg (Eds.), *The Working Alliance: Theory, Research and Practice* (pp. 38-50). New York: John Wiley & Sons.
- Maes, P. (1989). How to do the right thing. *Connection Science Journal*, 1(3).
- Malinowski, B. (1923). The problem of meaning in primitive languages. In C. K. Ogden & I. A. Richards (Eds.), *The Meaning of Meaning*: Routledge & Kegan Paul.
- Mallinckrodt, B. (1003). Session Impact, Working Alliance, and Treatment Outcome in Brief Counseling. *Journal of Counseling Psychology*, 40(1), 25-32.
- Mallinckrodt, B., & Nelson, M. L. (1991). Counselor Training Level and the Formation of the Psychotherapeutic Working Alliance. *Journal of Counseling Psychology*, 38(2), 133-138.
- Marcus, B. H., Selby, V. C., Niaura, R. S., & Rossi, J. S. (1992). Self-efficacy and the stages of exercise behavior change. *Research Quarterly for Exercise and Sport*, 63, 60-66.
- Marcus, B., & Owen, N. (1992). Motivational Readiness, Self-Efficacy and Decision-Making for Exercise. *J. of Applied Social Psychology*, 22(1), 3-16.
- Marsh, S. (1994). *Formalising Trust as a Computational Concept*. University of Stirling, Stirling, Scotland.
- McGuire, A. (1994). Helping Behaviors in the Natural Environment: Dimensions and Correlates of Helping. *Personality and Social Psychology Bulletin*, 20(1), 45-56.
- McNeill, D. (1992). *Hand and Mind: What Gestures Reveal about Thought*. Cambridge: Cambridge University Press.
- Mehrabian, A. (1972). *Nonverbal Communication*. New York: Aldine-Ahterton.
- Miell, D. (1987). Remembering relationship development: constructing a context for interactions. In R. Burnett & P. McGhee & D. Clarke (Eds.), *Accounting for Relationships* (pp. 60-73). New York: Methuen.
- Miller, C., & Larson, R. (1992). An explanatory and "argumentative" interface for a model-based diagnostic system. Paper presented at the Proceedings of the ACM Symposium on User Interface Software and Technology, Monterey, CA.
- Miller, G. A., Beckwith, R., Fellbaum, C., Gross, D., & Miller, K. (1993). *Introduction to Wordnet: An On-line Lexical Database*.

- Moon, Y. (1998). Intimate self-disclosure exchanges: Using computers to build reciprocal relationships with consumers (Working paper 99-059). Cambridge, MA: Harvard Business School.
- Morkes, J., Kernal, H., & Nass, C. (1998). Humor in Task-Oriented Computer-Mediated Communication and Human-Computer Interaction. Paper presented at the CHI 98.
- Mulken, S. v., Andre, E., & Muller, J. (1999). An Empirical Study on the Trustworthiness of Life-Like Interface Agents. Saarbrucken, Germany: DFKI.
- Nigg, C., Courneya, K., & Estabrooks, P. (1997). Maintaining Attendance at a Fitness Center: An Application of the Decision Balance Sheet. *Behavioral Medicine*, 23, 130-137.
- Nigg, C., Rossi, J., Norman, G., & Benisovich, S. (1998). Structure of decisional balance for exercise adoption. *Annals of Behavioral Medicine*, 20(S211).
- Okun, B. (1997). *Effective Helping: Interviewing and Counseling Techniques*. Pacific Grove, CA: Brooks/Cole.
- Okun, B. F. (1997). *Effective Helping: Interviewing and Counseling Techniques* (5 ed.). Pacific Grove, CA: Brooks/Cole Publishing.
- Pate, R. R., Pratt, M., Blair, S. N., Haskell, W. L., Macera, C. A., Bouchard, C., Buchner, D., Ettinger, W., Heath, G. W., King, A. C., Kriska, A., Leon, A. S., Marcus, B. H., Morris, J., Paffenbarger, R. S., Patrick, K., Pollock, M. L., Rippe, J. M., Sallis, J., & Wilmore, J. H. (1995). Physical Activity and Public Health: A Recommendation From the Centers for Disease Control and Prevention and the American College of Sports Medicine. *Journal of the American Medical Association*, 273(5), 402-407.
- Pautler, D. (1998). A Computational Model of Social Perlocutions. Paper presented at the COLING/ACL.
- Pelachaud, C., Badler, N., & Steedman, M. (1994). Generating Facial Expressions for Speech. *Cognitive Science*, 1994.
- Petty, R., & Wegener, D. (1998). Attitude Change: Multiple Roles for Persuasion Variables. In D. Gilbert & S. Fiske & G. Lindzey (Eds.), *The Handbook of Social Psychology* (pp. 323-390). New York: McGraw-Hill.
- Picard, R. (1997). *Affective Computing*. Cambridge, MA: MIT Press.
- Picard, R. W., Vyzas, E., & Healey, J. (2001). Toward Machine Emotional Intelligence: Analysis of Affective Physiological State. *IEEE Transactions Pattern Analysis and Machine Intelligence*, 23(10).
- Picard, R., & Klein, J. (2002). Computers that recognize and respond to user emotion: theoretical and practical implications. *Interacting with Computers*, 14, 141-169.
- Picard, R., & Scheirer, J. (2001, August). The Galvactivator: A Glove that Senses and Communicates Skin Conductivity. Paper presented at the International Conference on Human-Computer Interaction, New Orleans.

- Pierrehumbert, J., & Hirschberg, J. (1990). The Meaning of Intonational Contours in Interpretation of Discourse. In P. R. Cohen & J. Morgan & M. E. Pollack (Eds.), *Intentions in Communication* (pp. 271-311). Cambridge, MA: MIT Press.
- Pinto, B., Cherico, N., Szymanski, L., & Marcus, B. (1998). Longitudinal Changes in College Students' Exercise Participation. *College Health*, 47, 23-27.
- Planalp, S. (1993). Friends' and Acquaintances' Conversations II: Coded Differences. *Journal of Social and Personal Relationships*, 10, 339-354.
- Planalp, S. (1998). Communicating Emotion in Everyday Life: Cues, Channels and Processes. In P. Andersen & L. Guerrero (Eds.), *Handbook of Communication and Emotion*. New York: Academic Press.
- Planalp, S., & Benson, A. (1992). Friends' and Acquaintances' Conversations I: Perceived Differences. *Journal of Social and Personal Relationships*, 9, 483-506.
- Pollack, M. E. (1990). Plans as Complex Mental Attitudes. In P. R. Cohen & J. Morgan & M. E. Pollack (Eds.), *Intentions in Communication* (pp. 77-102). Cambridge, MA: MIT Press.
- Prevost, S., & Steedman, M. (1994). Specifying intonation from context for speech synthesis. *Speech Communication*, 15, 139-153.
- Prince, E. P. (1981). Toward a Taxonomy of Given-New Information. In Cole (Ed.), *Radical Pragmatics* (pp. 223-255): Academic Press.
- Prochaska, J. (2001). Helping Populations Progress Through Stages of Change, Talk given at Boston University. webcast <http://www.bu.edu/cpr/webcast/change.html>.
- Prochaska, J. (2003). Personal communication.
- Prochaska, J. O., Norcross, J. C., & Diclemente, C. C. (1994). *Changing for Good: A Revolutionary Six-Stage Program for Overcoming Bad Habits and Moving Your Life Positively Forward*. New York: Avon Books.
- Prochaska, J., & Marcus, B. (1994). The Transtheoretical Model: Applications to Exercise. In R. Dishman (Ed.), *Advances in Exercise Adherence* (pp. 161-180). Champaign, IL: Human Kinetics.
- Raue, P., & Goldfried, M. (1994). The Therapeutic Alliance in Cognitive-Behavior Therapy. In A. Horvath & L. Greenberg (Eds.), *The Working Alliance: Theory, Research and Practice* (pp. 131-152). New York: John Wiley & Sons.
- Reeves, B., & Nass, C. (1996). *The Media Equation*. Cambridge: Cambridge University Press.
- Reynolds, C., & Picard, R. W. (2001, August). Designing for Affective Interactions. Paper presented at the 9th International Conference on Human-Computer Interaction.
- Rich, C., & Sidner, C. L. (1998). COLLAGEN: A collaboration manager for software interface agents (TR-97-21a). Cambridge, MA: Mitsubishi Electric Information Technology Center America.
- Rich, C., Sidner, C. L., & Lesh, N. (2001). COLLAGEN: Applying Collaborative Discourse Theory to Human-Computer Interaction. *AI Magazine*.

- Richmond, V., & McCroskey, J. (1995). *Immediacy, Nonverbal Behavior in Interpersonal Relations* (pp. 195-217). Boston: Allyn & Bacon.
- Rickel, J., & Johnson, W. L. (1998). *Animated Agents for Procedural Training in Virtual Reality: Perception, Cognition and Motor Control*. *Applied Artificial Intelligence*.
- Riva, A., Smigelski, C., & Friedman, R. (2000). *WebDietAID: An Interactive Web-Based Nutritional Counselor*. Paper presented at the AMIA.
- Rose, C., Bodenheimer, B., & Cohen, M. (1998). Verbs and Adverbs: Multidimensional motion interpolation using radial basis functions. *IEEE CGAA(?)*(Fall).
- Sacks, H. (1995). *Lectures on Conversation*. Oxford: Blackwell.
- Safran, J., & Wallner, L. (1991). The relative predictive validity of two therapeutic alliance measures in cognitive therapy. *Psychological Assessment*, 3(2), 188-195.
- Sallis, J. (1997). Seven-Day Physical Activity Recall. *Medicine & Science in Sports & Exercise, Supplement*, S89-S103.
- Salovey, P., & Mayer, J. D. (1990). Emotional Intelligence. *Imagination, Cognition and Personality*, 9(3), 185-211.
- Scherer, K. (1981). Speech and emotional states. In J. Darby (Ed.), *Speech Evaluation in Psychiatry* (pp. 189-220): Grune and Stratton.
- Scherer, K., Ladd, D., & Silverman, K. (1984). Vocal cues to speaker affect: testing two models. *Journal Acoustical Society of America*, 76(5), 1346-1355.
- Schiffrin, D. (1987). *Discourse markers*. Cambridge: Cambridge University Press.
- Schneider, K. P. (1987). Topic selection in phatic communication. *Multilingua*, 6(3), 247-256.
- Schneider, K. P. (1988). *Small Talk: Analysing Phatic Discourse*. Marburg: Hitzeroth.
- Searle, J. (1969). *Speech Acts: An essay in the philosophy of language*: Cambridge University Press.
- Sidner, C. (1994). *An Artificial Discourse Language for Collaborative Negotiation*. Paper presented at the AAAI.
- Sidner, C. (2002, June 28-29). *Engagement between Humans and Robots for Hosting Activities*. Paper presented at the International CLASS Workshop on Natural, Intelligent and Effective Interaction in Multimodal Dialogue Systems, Copenhagen.
- Smith, J. (2000). *GrandChair: Conversational Collection of Family Stories*. MIT, Cambridge, MA.
- Spencer-Oatey, H. (1996). Reconsidering power and distance. *Journal of Pragmatics*, 26, 1-24.
- Sproull, L., Subramani, M., Kiesler, S., Walker, J., & Waters, K. (1997). When the Interface is a FAcE. In B. Friedman (Ed.), *Human Values and the Design of Computer Technology* (pp. 163-190). Stanford, CA: CSLI Publications.

- Stafford, L., & Canary, D. (1991). Maintenance Strategies and Romantic Relationship Type, Gender and Relational Characteristics. *Journal of Social and Personal Relationships*, 8, 217-242.
- Stone, M., & Doran, C. (1997). Sentence Planning as Description Using Tree-Adjoining Grammar. Paper presented at the ACL.
- Sunde, M., & Sandra, S. (1999). *Behavior Change in the Human Services: An Introduction to Principles and Applications* (4th ed.). London: Sage.
- Svennevig, J. (1999). *Getting Acquainted in Conversation*. Philadelphia: John Benjamins.
- Takeuchi, A., & Naito, T. (1995). Situated Facial Displays: Towards Social Interaction. Paper presented at the CHI '95.
- Tannen, D. (1993a). Introduction (Framing in Discourse). In D. Tannen (Ed.), *Framing in Discourse* (pp. 3-13). New York: Oxford University Press.
- Tannen, D. (1993b). What's in a Frame? Surface Evidence for Underlying Expectations. In D. Tannen (Ed.), *Framing in Discourse* (pp. 14-56). New York: Oxford University Press.
- Taylor, P., Black, A., & Caley, R. (1998). The Architecture of the Festival Speech Synthesis System. Paper presented at the 3rd ESCA Workshop on Speech Synthesis, Jenolan Caves, Australia.
- Thomason, R. (1990). Accommodation, Meaning, and Implicature: Interdisciplinary Foundations for Pragmatics. In P. Cohen & J. Morgan & M. Pollack (Eds.), *Intentions in Communication* (pp. 325-364). Cambridge, MA: MIT Press.
- Thorisson, K. R. (1997). Gandalf: An Embodied Humanoid Capable of Real-Time Multimodal Dialogue with People. Paper presented at the Autonomous Agents '97.
- Torres, O. E., Cassell, J., & Prevost, S. (1997, July 1997). Modeling Gaze Behavior as a Function of Discourse Structure. Paper presented at the First International Workshop on Human-Computer Conversation.
- Tracy, K., & Coupland, N. (1991). Multiple goals in discourse: An overview of issues. In K. Tracy & N. Coupland (Eds.), *Multiple goals in discourse* (pp. 1-13). Clevedon: Multilingual Matters.
- Traum, D., & Allen, J. (1994). Discourse Obligations in Dialogue Processing. Paper presented at the ACL '94.
- Tseng, S., & Fogg, B. J. (1999). Credibility and Computing Technology. *CACM*, 42(5), 39-44.
- Tudor-Locke, C. Pedometers are practical and accurate tools for measurement and motivation for physical activity.
- Turkle, S. (1995). *Life on the Screen: Identity in the Age of the Internet*. Englewood Cliffs, NJ: Simon & Schuster.
- Turkle, S. (2002). A new kind of object: From Rorschach to Relationship. *Edge: The Third Culture*. http://www.edge.org/3rd_culture/story/101.html

- Velicer, W., & Prochaska, J. (1999). An Expert System Intervention for Smoking Cessation. *Patient Education and Counseling*, 36, 119-129.
- Walker, M. A., Cahn, J. E., & Whittaker, S. J. (1997). Improvising linguistic style: Social and Affective Bases for Agent Personality. Paper presented at the Autonomous Agents 97, Marina Del Rey, CA.
- Wheless, L., & Grotz, J. (1977). The Measurement of Trust and Its Relationship to Self-Disclosure. *Human Communication Research*, 3(3), 250-257.
- Whittaker, S., & O'Conaill, B. (1997). The Role of Vision in Face-to-Face and Mediated Communication. In K. Finn & A. Sellen & S. Wilbur (Eds.), *Video-Mediated Communication* (pp. 23-49): Lawrence Erlbaum Associates, Inc.
- Williams, C., & Stevens, K. (1972). Emotions and speech: some acoustical correlates. *Journal Acoustical Society of America*, 52(4), 1238-1250.
- Winnicott, D. (1982). *Playing and Reality*. New York: Routledge.
- Wiseman, J. (1986). Friendship: Bonds and binds in a voluntary relationship. *Journal of Social and Personal Relationships*, 3, 191-211.
- Wish, M., Deutsch, M., & Kaplan, S. J. (1976). Perceived Dimensions of Interpersonal Relations. *Journal of Personality and Social Psychology*, 33(4), 409-420.
- Woods, W. A. (1986). Transition Network Grammars for Natural Language Analysis. In B. J. Grosz & K. S. Jones & B. L. Webber (Eds.), *Readings in Natural Language Processing* (pp. 71-88). Los Altos, CA: Morgan Kaufmann Publishers, Inc.
- Yan, H. (1999). *Semantic Features Selection in Speech & Gesture Generation* (Discourse Final Project).
- Yan, H. (2000). *Paired Speech and Gesture Generation in Embodied Conversational Agents*. MIT, Cambridge, MA.
- Yan, H., & Selker, T. (2000). Context-Aware Office Assistant. Paper presented at the International Conference on Intelligent User Interfaces, New Orleans.
- Young, R. M., & Moore, J. D. (1994). DPOCL: A Principled Approach to Discourse Planning. Paper presented at the Proceedings of the 7th International Workshop on Natural Language Generation, Kinnebunkport, ME.

Appendix A: Experimental Protocol for the Small Talk Experiment

Subject recruiting advertisement



TOUR THE MEDIA LAB

We need some folks to help us out with a cool experiment. It takes less than an hour, and you get ten bucks plus a tour of the lab.



You must have been born in the U.S. of A., and be a native speaker of English. Email us at GNL-EXPT@media.mit.edu to schedule a

\$10 for Cool Media Lab expert interest gnl-expt@media.mit.edu

Experimenter's Script

EMBODIED Condition

Welcome/Consent

[seat the subject experiment room]

Thanks for helping out with this experiment. Before we get started I need you to sign some consent forms.

[hand subject forms]

You can go ahead and read the consent forms now.

[Give subject time to read and sign consent forms.]

[Collect consent forms.]

Explain Task

Now I'm going to explain more about what you'll be doing. We're evaluating a life-sized, animated computer character named Rea. Rea plays the role of a real estate agent. She will ask you questions about the kind of apartment you would be interested in renting the next time you move, and she can show you some apartments that she has for rent. You are to play the role of someone looking for an apartment in the Boston area.

Is this clear?

You will interact with Rea the same way that you interact with another person, just by standing in front of her and talking to her.

Are you ready to continue?

Preparation

I'm now going to show you a short movie of Rea in which she is going to give you some additional instructions. This is just a recording; you will be meeting Rea after the movie is over.

[Turn monitor on.]

[Play intruct.mov on boa.]

[Stand just in front of the screen as if you were talking to Rea.]

OK. Rea is going to appear on this big screen. I need you to stand up to talk to her just like you would to another person.

[Wait for subject to stand up.]

[Maximize animator window on polong and switch display to projector.]

[Turn room lights off.]

In a moment, Rea is going to begin talking to you about your housing needs. I'll wait outside. When she finishes and says goodbye knock on the door and I'll come back in.

[Leave room and walk to wizard room quietly.]

[Start video taping.]

[Make Rea face user and do greeting.]

<INTERACTION>

[Leave wizard room.]

[If subject has knocked on door, enter, otherwise wait 30 seconds, then enter.]

All finished?

[Turn room lights on, switch display to polong monitor, turn monitor off.]

Completing the questionnaire

You can have a seat now. I'd now like you to complete two short questionnaires, this is the first part. I'll wait outside. Just knock on the door when you're finished.

[Hand subject questionnaire RST-1.]

[When first part is complete, collect RST-1 and give subject RST-2.]

I'll wait outside again. Just knock on the door when you're finished.

[Collect the questionnaire]

Debrief

[Start audio tape.]

First, let me ask you what your overall impression of the experiment was. Did anything seem strange to you?

[Discuss subject's reaction.]

How natural did you feel the interaction was?

[Discuss subject's reaction.]

What did you think of Rea's social small talk and her real estate talk?

[Discuss subject's reaction.]

Did you think you were interacting with a computer or a person?

[Discuss subject's reaction.]

I'd like to take a moment to go over the purpose of the study, and tell you a little bit about what we're investigating. We are interested in whether social dialog, also known as "small talk" or "chit chat", affects your perception of the credibility and trustworthiness of a computer agent. The study has two conditions. In one condition we had the agent use a human model of small talk, so that it reproduced it in as natural a manner as possible. In the second condition, we had the agent use no small talk at all.

[Tell subject what condition they were in. Discuss their reaction.]

It was very important to us that problems with the speech recognition technology not effect the things we were interested in measuring. For example, if you had an accent and Rea could not understand you very well, you might have thought that she is rude or stupid. Although we are working on a system that actually responds to speech input, to ensure that these problems not effect our results we had a human observer listening to you and controlling Rea's responses. Thank you for your help, and I apologize for any discomfort this mild deception might have caused you. We hope the experience was an enjoyable and interesting one for you.

I'd just like to add one more thing before you go. It's really important that you not discuss this experiment with anyone else, because if someone comes and does the experiment knowing what we are trying to study they can throw the results off.

[Discuss subject's reaction.]

Finally, I need you to fill out this form in order to get paid.

[Give subject payment voucher to fill out.]

If you follow me, we just need to make a copy of this for our records.

[Lead subject to Andrew's desk. Photocopy payment voucher.]

Thanks again for your help. You can collect payment for this experiment by taking this form to the cashier's office in the infinite corridor, room 10-180.

[Give subject copy of payment voucher.]

Consent Form

You are being invited to participate in a research study. This form is designed to provide you with information about this study. The Principal Investigator or representative will describe this study to you and answer any of your questions. If you have any questions or complaints about the informed consent process or the research study, please contact the MIT Committee on Use of Human Experimental Subjects (COUHES), the committee that protects human subjects, at 253-6787.

Your participation in the following experiment is completely voluntary. You are free to withdraw this consent at any time, for any reason, and to request that any data collected be destroyed. If at any time you feel uncomfortable, or unsure that you wish your results to be part of the experiment, you may discontinue your participation with no repercussions.

The purpose of this study is to understand the role of technology in some aspects of human communication. To this end, the experimenters are analyzing people having conversations using a computer system.

The full experiment is expected to take one hour or less, and you will be compensated with \$10, whether you complete all the tasks or not.

In order to explore the results, we will videotape your conversation for subsequent transcription and study.

After completing the conversational task, you will be asked to fill out a questionnaire about your reactions to the task. When the experiment is over, the investigator will fully debrief you about the goals of the study and answer any questions you might have regarding the study and the planned use of the videotapes.

All specific information divulged in this experiment will be kept confidential by the researchers. Your participation will be videotaped and only the researchers will view your tapes, unless you specify otherwise. You can request to view your videotape after the experiment, and can withdraw your videotape from the study if requested within four weeks of the date of the experiment. Note that all videotapes will be stored in a locked cabinet, accessible only to the principal investigators. In the unlikely event that it becomes impossible for the principal investigators to provide such a secure storage space, the videotapes will be destroyed. In addition, the videotape and your responses will be completely anonymous. All data will be associated only with an ID number.

Please read the following and sign on the lines below:

"I, the undersigned, have read and understood the explanations of the following research project and voluntarily consent to my participation in it. I understand that my responses will remain confidential and that I may terminate my participation at any time.

In the unlikely event of physical injury resulting from my participation in this research, I understand that medical treatment will be available from the MIT Medical Department, including first aid emergency treatment and follow-up care as needed, and that my insurance carrier may be billed for the cost of such treatment. However, no compensation can be provided for medical care apart from the foregoing. I further understand that making such medical treatment available or providing it, does not imply that such injury is the investigator's fault. I also understand that by my participation in this study I am not waving any of my legal rights.

I understand that I may also contact the Chairman of the Committee on the Use of Humans as Experimental Subjects, MIT 253-6787, if I feel that I have been treated unfairly as a subject.

Name _____

Signature _____ Date _____

Use of videotape

We would also appreciate your consent to use very short clips of the videotapes to show in educational contexts. If we do use short clips of you, we will edit them to only show short examples of your interaction with the system.

This consent is entirely separate from your consent to participate in the experiment and may be withdrawn at any time in the future. To give your permission to have an excerpt from the videotape potentially shown for research and teaching purposes, please sign below. Please note that if you do not wish to give your permission for this, you may still participate in the study.

I have read and understand the above and agree to potentially have an excerpt from my videotaped interaction shown for educational purposes.

Name _____

Signature _____ Date _____

Questionnaire Part 1

How much monthly rent would you be willing to pay for the second apartment Rea showed you?

\$

How much did you like Rea?

not at all • • • • • • • • • • very much

Would you enjoy working with Rea again?

not at all • • • • • • • • • • very much

How well did Rea understand you?

not at all • • • • • • • • • • perfectly

How well do you feel Rea knows you and your apartment needs?

not at all • • • • • • • • • • very well

How well do you feel you know Rea?

not at all • • • • • • • • • • very well

How would you characterize your relationship with Rea?

complete stranger • • • • • • • • • • close friend

To what extent do the following words describe your interaction with Rea? Note that you are NOT evaluating Rea, but the *interaction* between you and Rea.

	Describes very poorly										Describes very well									
Awkward	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Comfortable	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Efficient	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Engaging	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Enjoyable	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Familiar	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Fun	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Interesting	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Involving	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Natural	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Satisfying	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Smooth	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Successful	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•
Tedious	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•	•

To what extent do the following words describe Rea? Note that you are evaluating *Rea* now, NOT the interaction.

	Describes very poorly					Describes very well				
Competent	•	•	•	•	•	•	•	•	•	•
Credible	•	•	•	•	•	•	•	•	•	•
Expert	•	•	•	•	•	•	•	•	•	•
Friendly	•	•	•	•	•	•	•	•	•	•
Informed	•	•	•	•	•	•	•	•	•	•
Intelligent	•	•	•	•	•	•	•	•	•	•
Knowledgeable	•	•	•	•	•	•	•	•	•	•
Lifelike	•	•	•	•	•	•	•	•	•	•
Likeable	•	•	•	•	•	•	•	•	•	•
Pleasant	•	•	•	•	•	•	•	•	•	•
Reliable	•	•	•	•	•	•	•	•	•	•
Warm	•	•	•	•	•	•	•	•	•	•

Please state your opinion of Rea on each of the scales below.

benevolent	•	•	•	•	•	•	•	•	exploitive
trustworthy	•	•	•	•	•	•	•	•	untrustworthy
confidential	•	•	•	•	•	•	•	•	divulging
safe	•	•	•	•	•	•	•	•	dangerous
candid	•	•	•	•	•	•	•	•	deceptive
not deceitful	•	•	•	•	•	•	•	•	deceitful
trustful of Rea	•	•	•	•	•	•	•	•	distrustful of Rea
straightforward	•	•	•	•	•	•	•	•	tricky
respectful	•	•	•	•	•	•	•	•	disrespectful
considerate	•	•	•	•	•	•	•	•	inconsiderate
honest	•	•	•	•	•	•	•	•	dishonest
reliable	•	•	•	•	•	•	•	•	unreliable
faithful	•	•	•	•	•	•	•	•	unfaithful
sincere	•	•	•	•	•	•	•	•	insincere
careful	•	•	•	•	•	•	•	•	careless

Questionnaire Part 2

Please check the item in each question that best describes you.

Are you more inclined to be:

easy to approach

somewhat reserved

In company, do you:

initiate conversation

wait to be approached

Does new and nonroutine interaction with others:

stimulate and energize you

tax your reserves

Do you prefer:

many friends with brief contact

a few friends with more lengthy contact

In social groups do you:

keep abreast of others' happenings

get behind on the news

Do you:

speak easily and at length with strangers

find little to say to strangers

**To what extent do the following words describe you?
Circle one dot on each line.**

bashful

not at all • • • • • • • • • • perfectly

cheerful

not at all • • • • • • • • • • perfectly

enthusiastic

not at all • • • • • • • • • • perfectly

extroverted

not at all • • • • • • • • • • perfectly

introverted

not at all • • • • • • • • • • perfectly

inward

not at all • • • • • • • • • • perfectly

jovial

not at all • • • • • • • • • • perfectly

outgoing

not at all • • • • • • • • • • perfectly

perky

not at all • • • • • • • • • • perfectly

shy

not at all • • • • • • • • • • perfectly

undemonstrative

not at all • • • • • • • • • • perfectly

unrevealing

not at all • • • • • • • • • • perfectly

unsparkling

not at all • • • • • • • • • • perfectly

Final Page

Did you think the second apartment that Rea showed you was worth the asking price?

definitely not • • • • • • • • • • definitely worth it

How do you feel about the amount of social "chit chat" Rea did?

not enough • • • • • • • • • • too much

How much did Rea's social "chit chat" get in the way of your business interaction?

not at all • • • • • • • • • • very much

How much did Rea's social "chit chat" help your business interaction?

not at all • • • • • • • • • • very much

How do you feel about how quickly Rea "got down to business"?

too long • • • • • • • • • • too quickly

REA Script for SOCIAL Condition

Note: Words in '{...}' braces are commands to the animation system. Words in '<...>' brackets are values entered by the Wizard.

1. Nice to meet you.
2. How are you doing?
3. How do you like being on camera? I'm a ham, but I know it makes some people nervous.
4. Well. This whole set-up can sometimes feel a little strange; I hope you're more or less comfortable.
5. How {beat} about this weather? Aren't summers in Boston awful? The air is so thick you can cut it with {beat} a knife.
6. I think autumn is the most beautiful season in New England. {beat} Not to mention that that's when the new
7. students arrive, and I get to meet lots of interesting people.
8. So, Have you been in the Media Lab before?
9. It's pretty cool isn't it?
10. How do you like Boston?
11. Boston is certainly more expensive than it used to be.
12. Anyway, I have {beat} several apartments in my database that you might be interested in.
13. {armshape arm_round4} Where would you like to live?
14. Do you need access to the subway?
15. How many {beat} bedrooms do you need?
16. Is one bath enough?
17. That reminds me. Last year I rented an apartment to an MIT student that actually had a {beat} hot tub in the

18. {beat} bathroom of a studio apartment. It was {beat} unbelievable. Great for parties, though.
19. So, do you need any outdoor space?
20. Let's talk about {beat} style. Do you like {beat} older victorian style, with hardwood floors, or {beat} more modern?
21. Do you prefer {beat} large closets, or are {beat} small closets OK if you have some storage in a basement?
22. Would you prefer to rent in {beat} part of a house or an {beat} apartment building?
23. It's amazing how many houses have been turned into rental units in Boston. It's getting almost impossible
24. to find single family homes for sale anymore.
25. {beat} Anyway. What can you afford for monthly rent?
26. Let me see.
27. {slide 7} I have an {armshape arm_round4} apartment in a house in <CITY> and it just became available last week.
28. It's a little larger than you need, but it is very beautiful.
29. The flat has <BEDROOMSX2> and two baths.
30. It's not too far from MIT.
31. Let me show you inside.
32. {slide 9} {armshape arm_user_deictic} This is the living room. It is very spacious.
33. {slide 10} {armshape arm_user_deictic} The master bedroom is upstairs. The bathroom is next to it.
34. {slide 8} {armshape arm_user_deictic} This is the kitchen. There is an antique stove here.
35. {slide 9} What do you think?
36. The owner is asking <RENTX2> a month, firm.

37. You know, I rented an apartment last month in downtown Boston for \$4,200 a month.
You'd think the
38. floors were gold-plated or something.
39. Oh, I do have one more property that I'd like to show you. Is that OK?
40. {slide 2} I think you'll like it. It's in a building in <CITY>.
41. It's five minutes to the T station.
42. {slide 4} {armshape arm_user_deictic} This is the living room. The apartment has
<BEDROOMS> and one bath.
43. {slide 3} {armshape arm_user_deictic} It has a modern kitchen complete with a
dishwasher.
44. {slide 5} {armshape arm_user_deictic} This is the master bedroom. It is very
spacious and has wall-to-wall carpeting.
45. {slide 4} Would this work for you?
46. The landlord is asking <RENTX1.5> per month, but I think he's negotiable.
47. {slide 1} Why don't you think about these two apartments and give me a call if you
want to talk more?
48. Well, I have to run to another appointment now. Why don't you give me a call if
you're interested
49. in one of these apartments.
50. It was very nice talking to you. I hope
51. that this was helpful, and that we can work together some more so that I can help you
find the exact
52. apartment you need.

REA Script for TASK Condition

1. I have {beat} several apartments in my database that you might be interested in.
2. So, {armshape arm_round4} Where would you like to live?
3. Do you need access to the subway?
4. How many {beat} bedrooms do you need?
5. Is one bath enough?
6. Do you need any outdoor space?
7. Let's talk about {beat} style. Do you like {beat} older victorian style, with hardwood floors, or {beat} more modern?
8. Do you prefer {beat} large closets, or are {beat} small closets OK if you have some storage in a basement?
9. Would you prefer to rent in {beat} part of a house or an {beat} apartment building?
10. What can you afford for monthly rent?
11. Let me see.
12. {slide 7} I have an apartment in a house in <CITY> and it just became available last week.
13. It's a little larger than you need, but it is very beautiful.
14. The flat has <BEDROOMSX2> and two baths.
15. It's not too far from MIT.
16. Let me show you inside.
17. {slide 9} {armshape arm_user_deictic} This is the living room. It is very spacious.
18. {slide 10} {armshape arm_user_deictic} The master bedroom is upstairs. The bathroom is next to it.

19. {slide 8} {armshape arm_user_deictic} This is the kitchen. There is an antique stove here.
20. {slide 9} What do you think?
21. The owner is asking <RENTX2> a month, firm.
22. Oh, I have one more property that I'd like to show you. Is that OK?
23. {slide 2} I think you'll like it. It's in a building in <CITY>.
24. It's five minutes to the T station.
25. {slide 4} {armshape arm_user_deictic} This is the living room. The apartment has <BEDROOMS> and one bath.
26. {slide 3} {armshape arm_user_deictic} It has a modern kitchen complete with a dishwasher.
27. {slide 5} {armshape arm_user_deictic} This is the master bedroom. It is very spacious and has wall-to-wall carpeting.
28. {slide 4} Would this work for you?
29. The landlord is asking <RENTX1.5> per month, but I think he's negotiable.
30. {slide 1} Why don't you think about these two apartments and give me a call if you want to talk more?
31. Well, I have to run to another appointment now. Why don't you give me a call if you're interested
32. in one of these apartments.
33. It was very nice talking to you. I hope
34. that this was helpful, and that we can work together some more so that I can help you find the exact
35. apartment you need.

Appendix B: Protocol for the Exercise Advisor Experiment

Subject Recruiting Advertisements

Newspaper Ad

Want \$50?



Volunteers needed for a web-based study of physical activity among MIT students and staff. It takes only 10-15 minutes a day for four weeks (plus time for any additional exercise we can persuade you to do), and you can do it in your bathrobe (assuming you have a home computer with an internet connection).

Participants will receive \$25 in cash and \$25 worth of fitness goodies.

You must be in generally good health and interested in increasing your level of physical activity.

To see if you qualify, check out:

<http://exercise.media.mit.edu:2002/>

Recruitment Posters

Live Long and Prosper
 Maintaining a minimum level of daily physical activity can increase your life expectancy by two years.



Volunteers needed for a web-based study of physical activity among MIT students and staff. It takes only 10-15 minutes a day for four weeks (plus any additional exercise we can persuade you to do), and you can do it in your bathrobe (assuming you have a home computer with an internet connection).

Participants will receive \$25 in cash and \$25 worth of fitness goodies. You must be in generally good health and interested in increasing your level of physical activity.

Beam up to <http://exercise.media.mit.edu:2002/> to see if you qualify.

Physical Activity Experiment http://exercise.media.mit.edu:2002/											
---	---	---	---	---	---	---	---	---	---	---	---



Want to Get in Shape?

We'll pay you to get started on a new exercise program.

Volunteers needed to try out a new exercise program. It takes only 10-15 minutes a day on your home computer for four weeks (plus time for any additional exercise we can persuade you to do).

Participants will receive \$25 in cash and \$25 worth of fitness goodies. You must be in generally good health and interested in increasing your level of physical activity.

Check out <http://exercise.media.mit.edu:2002/> to see if you qualify.

Physical Activity Experiment http://exercise.media.mit.edu:2002/											
---	---	---	---	---	---	---	---	---	---	---	---

Need to kick-start your brain?

Exercise increases levels of norepinephrine and endorphin in your head, increases your self-confidence and improves your ability to deal with stress!



Volunteers needed for a web-based study of physical activity among MIT students and staff. It takes only 10-15 minutes a day for four weeks (plus any additional exercise we can persuade you to do), and you can do it in your bathrobe (assuming you have a home computer with an internet connection).

Participants will receive \$25 in cash and \$25 worth of fitness goodies. You must be in generally good health and interested in increasing your level of physical activity.

Go to <http://exercise.media.mit.edu:2002/> to see if you qualify.

Physical Activity Experiment http://exercise.media.mit.edu:2002/										
---	---	---	---	---	---	---	---	---	---	---



Want \$50?



Volunteers needed for a web-based study of physical activity among MIT students and staff. It takes only 10-15 minutes a day for four weeks (plus any additional exercise we can persuade you to do), and you can do it in your bathrobe (assuming you have a home computer with an internet connection).

Participants will receive \$25 in cash and \$25 worth of fitness goodies. You must be in generally good health and interested in increasing your level of physical activity.

Dial in to <http://exercise.media.mit.edu:2002/> to see if you qualify.

Physical Activity Experiment http://exercise.media.mit.edu:2002/										
---	---	---	---	---	---	---	---	---	---	---

Experimenter's Intake Script

1. Thanks for helping out with this experiment. First I'd like you to fill out a brief questionnaire, this is a copy of the questionnaire you filled out on the web. *[Give subject screening questionnaire, including contact information, NOTE: this form will have their CONDITION and SUBJECT ID pre-printed on it][Give subject time to fill out.][Collect questionnaire and verify that all participation requirements are met]*

[If Subject has any medical conditions indicated on the form:]

I'm sorry, but in order to participate in this study you must be able to increase your level of physical activity without any possibility of health risks. This study is designed for a general population, and you may require special considerations given your health status. I would recommend talking to your personal physician about a customized exercise program that would be safe for you. *[Dismiss subject]*

[If Subject indicates they have participated in a Media Lab experiment before, find out which one. If GNL or AFFECTIVE, or if subject is affiliated with Media Lab]

I'm sorry, but because of your [affiliation/past experience] with the Media Lab you may know too much about the goals of the study for us to be able to use you. *[Dismiss subject]*

2. Now I'm going to tell you about what you're going to be doing in the study. We're investigating exercise behavior and fitness in the college population. In order to measure how much physical activity you are doing each day we need you to wear this pedometer *[show pedometer]* for the next month. During this month, we need you to access a web site from home once a day and spend approximately 10 to 15 minutes filling out a form, recording your pedometer reading and describing how much physical activity you engaged in during the previous 24 hours. The very first time you access the site you will be asked to estimate this information for the prior week. In addition, two months from the start of the experiment we will contact you via email and ask you to take 15 minutes and fill out the one-week activity forms again. The web site will guide you through the necessary steps each day; to stay in the study all that is required is that you wear the pedometer and login each day for 10-15 minutes. If you go away for the weekend, that's OK, you can miss up to 3 days without any problems. Are you able to help us out with this?

You can only access the website from a customized browser program that we will give you. You should install this on your home computer before the experiment begins.

Each time you fill out the activity form, you will be presented with additional information about exercise and fitness, and will occasionally be asked to fill out additional web forms related to your activity level or to rate some aspect of the Fit Track system.

There is also an animated character on the web site who will talk with you about your exercise experience. The character talks to you using synthesized speech, so it is important that you have your computer speakers or headphones on when you access the site.

All of this is described on an instruction sheet I'm going to give you. Do you have any questions? *[Discuss any questions the subject has]*

3. You meet all of the inclusion criteria for my study, and I would like to offer you a spot. But before you decide to accept, I want to explain the position I'm in so you can decide whether or not to enter. Unless you are almost positive that you will be able to complete the whole two months, please, please, do not enter the study.

If you drop out without completing the study, I can't use ANY of your data. So I'll have to replace you with another subject, and I might not be able to find one who can finish in time and I only have a couple of extra pedometers. I desperately need to finish this study so I can graduate in the spring, which means I need virtually all of the subjects who initially enter the study to complete it.

If you aren't sure you can or really want to complete the study, please don't agree to start. If you don't complete the study, it will waste your time and mine, and will decrease my chances of graduating.

If you are just not sure you will be able to start and stick to an exercise program, but are almost certain you can try and will be able to complete the forms on days 1-30 and day 60, that's all I can possibly ask of you, so please sign up. If the program doesn't seem to be helping you, it's especially important that you keep entering the data. It's just as important for me to learn what doesn't work as what does work. If my study doesn't report techniques that are not effective, other researchers will keep trying these techniques. Obviously, if you get mono and end up in the infirmary for a few weeks, I will understand. That's why I got money to run a few extra subjects, in case unavoidable circumstances prevent a couple subjects from finishing.

The most important thing to me is that you complete the study, and the second most important thing to me is that hopefully the study enables you to start an exercise program, and that I learn more about how to help people start and maintain exercise programs. So if during the study you have any problems, questions, comments, or suggestions, please contact me right away. You won't be bothering me – I'd rather have you e-mail on the weekend or at night, than not – and you can miss up to 3 days if you go away and don't have access to the computer program and still complete enough forms to complete the study. If there's anything I can do to help you finish the study, I will do it.

One other important thing is that you not talk to people about the study until you are finished. We are experimenting with different techniques for tracking your exercise behavior, and if you talk to someone else in the study who is using a different set of techniques that can ruin the results.

4. Now I need you to sign some consent forms. *[hand subject forms]* You can go ahead and read the consent forms now. *[Give subject time to read and sign consent forms.] [Collect consent forms.]*

5. Now I'm going to show you how to use your pedometer. It's very simple, you just wear it on your belt or waistband close to the midline of your thigh *[demonstrate]*. It's important that you attach this safety strap to some article of clothing so that if you run into a desk or doorway the pedometer doesn't accidentally get knocked off without you knowing about it. At the start or end of each day, simply write down the pedometer reading--it's just the number of steps you've taken during the day--then push the reset button. Any questions? *[answer any questions] [hand subject pedometer]* This is yours to keep.

6. Here are instructions for installing and running the software. *[Hand subject instructions with their USERID]*. Your USERID is written on the top, and you will need it to access the web site. Also written here is the email address for technical support. It is very important for the study that if you ever have any problems with any aspects of the web site that you send an email to technical support immediately so that it can get fixed as soon as possible. Also, if at any time you suffer a physical injury from your activities please notify us immediately at the same address.

7. Do you have a high-speed Internet connection?

[IF YES] You can download the software from the URL on the instructions sheet.

[IF NO] Here is a CD with the software on it. *[Hand subject CD]*

Just follow the installation instructions on the page, if you have any problems send us an email. You should install the software as soon as possible to make sure it installs correctly. We will then send you an email when its time for you to start running the software.

8. Finally, I'd like you to fill out a few questionnaires on the computer. Remember, your responses on these are completely voluntary. When asked for your USER ID, use the number on the instruction sheet I gave you.

[Seat Subject at Computer, launch Intake questionnaire]

<http://cgi.media.mit.edu:8080/Intake>

[If crashes or other problem have them fill out paper forms.]

9. Thanks again for helping out. Just remember two things: 1) install the software as soon as you can and let us know if you have any problems; and 2) just sit tight--don't use the pedometers or run the software--until we send you an email to start, around October 14th, then you can just follow the instructions that I gave you.

Consent Form

You are being invited to participate in a research study. This form is designed to provide you with information about this study. The Principal Investigator or representative will describe this study to you and answer any of your questions. If you have any questions or complaints about the informed consent process or the research study, please contact the MIT Committee on Use of Human Experimental Subjects (COUHES), the committee that protects human subjects, at 253-6787.

Your participation in the following experiment is completely voluntary. You are free to withdraw this consent at any time, for any reason, and to request that any data collected be destroyed. If at any time you feel uncomfortable, or unsure that you wish your results to be part of the experiment, you may discontinue your participation with no repercussions.

The purpose of this study is to understand the role of technology in some aspects of human health behavior. To this end, the experimenters are analyzing people's behavior using computers to track their physical activity over a period of time.

The full experiment is expected to take approximately 10 minutes a day of your time over a one month duration. During this time you will also be asked to wear a small digital pedometer that records the number of steps you take each day. An additional 10 minutes of your time will be required two months after the start of the experiment. As way of compensation you will be given the pedometer to keep, whether you complete all the tasks or not. In addition, if you complete all tasks in the study you will be compensated with an additional \$25 in cash. Participation in this experiment should also increase your awareness of your physical activity level relative to guidelines published by several health organizations, and provide you with information about the benefits of exercise.

The daily interactions involve accessing a web site at which you will be asked to fill out forms describing the amount of physical activity you have engaged in and other information related to your exercise behavior. Your responses on all forms and questionnaires are always voluntary and you can always choose to not reply to a particular question if you are uncomfortable doing so.

When the experiment is over, the investigator will fully debrief you about the goals of the study and answer any questions you might have regarding the study and the planned use of the data.

All specific information divulged in this experiment will be kept confidential by the researchers. Your data will be recorded in a database indexed only by an anonymous ID number, and only the researchers will have access to it. You can request to view your data after the experiment, and can withdraw your data from the study if requested within four weeks of the end of the experiment. However, today you will be asked to provide us with

some demographic information about yourself. Given that, we cannot completely assure your anonymity in the experiment.

Please read the following and sign on the lines below:

"I, the undersigned, have read and understood the explanations of the following research project and voluntarily consent to my participation in it. I understand that my responses will remain confidential and that I may terminate my participation at any time.

In the unlikely event of physical injury resulting from my participation in this research, I understand that medical treatment will be available from the MIT Medical Department, including first aid emergency treatment and follow-up care as needed, and that my insurance carrier may be billed for the cost of such treatment. However, no compensation can be provided for medical care apart from the foregoing. I further understand that making such medical treatment available or providing it, does not imply that such injury is the investigator's fault. I also understand that by my participation in this study I am not waving any of my legal rights.

I understand that I may also contact the Chairman of the Committee on the Use of Humans as Experimental Subjects, MIT 253-6787, if I feel that I have been treated unfairly as a subject.

Name_____

Signature_____Date_____

Intake Screening Form

USERID: _____

NAME: _____

EMAIL: _____ **PHONE:** _____

Circle your response to each of the following:

1. Has your doctor ever said that you have a heart condition and that you should only do physical activity recommended by a doctor? YES/NO
2. Do you feel pain in your chest when you do physical activity? YES/NO
3. In the past month, have you had chest pain when you were not doing physical activity? YES/NO
4. Do you lose your balance because of dizziness or do you ever lose consciousness?
YES/NO
5. Do you have a bone or joint problem that could be made worse by a change in your physical activity? YES/NO
6. Is your doctor currently prescribing drugs (for example, water pills) for you blood pressure or heart condition? YES/NO
7. Do you know of any other reason why you should not do physical activity? YES/NO
8. Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week.

Question: Are you physically active according to that definition?

- a) Yes, I have been for MORE than 6 months;
 - b) Yes, I have been for LESS than 6 months;
 - c) No, but I intend to in the next 30 days;
 - d) No, but I intend to in the next 6 months;
 - e) No, and I do NOT intend to in the next 6 months.
9. Do you own a personal computer with the following minimum resources? YES/NO
 - Microsoft Windows XP, 2000, NT, or 98
 - Pentium II, 233 MHz; 64 MB RAM or better
 - Sound card and speakers or headphone
 - 20 MB free hard disk space
 10. Are you affiliated with the Media Lab or have you participated in any experiments at the Media Lab before? YES/NO



Thank you for volunteering to participate in this exercise study. Please take a few moments to fill out some background information.

Page 1 of 6

1. UserID:	<input type="text"/>
2. Email Address:	<input type="text"/>
3. Password:	<input type="text"/>
Confirm Password:	<input type="text"/>
4. First name (what you go by):	<input type="text"/>
5. Age:	<input type="text"/>
6. Gender:	<input type="radio"/> Male <input type="radio"/> Female
7. Ethnic Background:	<input type="text" value="Choose One"/>
8. Marital Status:	<input type="text" value="Choose One"/>
9. Highest Degree Obtained:	<input type="text" value="Choose One"/>
10. What is your occupation?	<input type="text" value="Choose One"/>

Page 2 of 6

1. Do you feel that you are in good health overall?	<input type="text" value="Choose One"/>
2. Are you concerned about improving your health?	<input type="radio"/> Yes <input type="radio"/> No
3. How would you describe your diet?	<input type="text" value="Choose One"/>
4. Are you concerned about improving your diet?	<input type="radio"/> Yes <input type="radio"/> No
5. How would you describe your sleep?	<input type="text" value="Choose One"/>
6. Are you concerned about improving your sleep?	<input type="radio"/> Yes <input type="radio"/> No
7. How would you describe your fitness level?	<input type="text" value="Choose One"/>
8. Are you concerned about improving your fitness level?	<input type="radio"/> Yes <input type="radio"/> No

Page 3 of 6

1. What is the reason that you want to increase your physical activity (check all that apply)?

- Overall health
- Physical appearance
- Lose Weight
- Fitness in and of itself
- Gain muscle mass
- Decrease anxiety and depression
- Increase ability to study effectively

2. Have you had a period in your life when you exercised more than you do now?

Yes No

If yes, for the most recent time period,

3. When did you start?

5. When did you end?

7. What did your exercise consist of?

8. Do you play any team or individual competitive sports on a regular basis, either intramural or intercollegiate?

Yes
 No

9. Height: feet inches

11. Weight: lbs

Page 4 of 6

To what extent would you agree with the use of the following words to describe you?

1. Bashful	<input type="text"/>
2. Cheerful	<input type="text"/>
3. Enthusiastic	<input type="text"/>
4. Extroverted	<input type="text"/>
5. Introverted	<input type="text"/>
6. Inward	<input type="text"/>
7. Jovial	<input type="text"/>
8. Outgoing	<input type="text"/>
9. Perky	<input type="text"/>
10. Shy	<input type="text"/>
11. Undemonstrative	<input type="text"/>
12. Unrevealing	<input type="text"/>
13. Unsparkling	<input type="text"/>
14. Vivacious	<input type="text"/>
15. Silent	<input type="text"/>
16. Unshy	<input type="text"/>

Page 5 of 6

To what extent would you agree with the use of the following words to describe you?

1. Domineering	<input type="text"/>
2. Meek	<input type="text"/>
3. Unaggressive	<input type="text"/>
4. Self-effacing	<input type="text"/>
5. Dominant	<input type="text"/>
6. Self-confident	<input type="text"/>
7. Forceful	<input type="text"/>
8. Self-doubting	<input type="text"/>
9. Unauthoritative	<input type="text"/>
10. Firm	<input type="text"/>
11. Self-assured	<input type="text"/>
12. Unbold	<input type="text"/>
13. Un-self-conscious	<input type="text"/>
14. Assertive	<input type="text"/>
15. Timid	<input type="text"/>
16. Forceless	<input type="text"/>

Page 6 of 6

To what extent would you agree with the use of the following words to describe you?

1. warm	<input type="text"/>
2. sympathetic	<input type="text"/>
3. disagreeable	<input type="text"/>
4. friendly	<input type="text"/>
5. cold	<input type="text"/>
6. hard-hearted	<input type="text"/>

Subject Take Home Instructions

1. Installing the FitTrack Software

REQUIREMENTS

The software takes approximately 30MB of hard disk space, and you should have speakers and at least a 800x600 pixel display. Otherwise, just about any windows box should work (it's been tested on 98, NT, 2000, and XP).

INCOMPATIBILITY WARNING: If you use the Microsoft Speech API or any product that uses the speech synthesizer part of this package (e.g., Microsoft Agent, Bonzai Buddy, etc.), this will overwrite the existing speech synthesizer installation causing these products to not work properly.

DOWNLOAD

If you did not receive a CD during your initial interview at the Media Lab, you can download the software installer from:

<http://cgi.media.mit.edu:8080/SetupFitTrack.exe>

INSTALLATION:

1. Make sure you have Administrator privileges on your machine (if applicable).
 2. Launch the SetupFitTrack.exe application.
 3. When it asks for a user ID enter the user ID written above.
 4. When it asks if you want to install 'Flash Active X' answer YES
- That's it.

2. Running the Software

Make sure your audio is turned up and your speakers or headphones are working. Simply launch the FitTrack application from your Start menu. When done, just exit the application by clicking the close box at the upper right-hand corner of the window.

If you try to run the software before we notify you to begin the study, you will get an error message when you try to launch the FitTrack application.

3. Using the Pedometer

Just place the pedometer on the waistband of your clothing or belt. Position it close to the midline of your thigh in an upright position. It should fit snug to your body and not hang loose. Be sure to use attach the safety strap to some article of clothing so that if the pedometer gets knocked off (e.g., going through a doorway) you won't lose it (a very common problem). We have very few spare pedometers, but if you do lose yours please contact us immediately. Avoid dropping or crushing the pedometer, exposing it to excessive moisture, or forcing the clip onto a belt (which can cause the clip to break).

Every day during the study: clear the pedometer first thing in the morning (by pushing the reset button), wear it all day, then write down the number of steps you walked before going to bed at night (the FitTrack software will ask you to enter your steps for the previous day each time you access it).

4. Questions? Problems?

Feel free to contact us at any time if you have any questions or are having any problems at FitTrack@media.mit.edu. There is also a feedback form--available to you every day through the software--that you can use to report problems, give us feedback, or just say hi.

REMINDERS:

- **PLEASE DON'T USE THE PEDOMETERS OR SOFTWARE UNTIL WE EMAIL YOU TO START.**
- **PLEASE DON'T TALK ABOUT THE STUDY UNTIL YOU HAVE BEEN PAID.**

Activity Log Form

■ Log your activity for **December 2, 2002**

Sleep (night before) **Went to bed at:**
7pm ▾
Got out of bed at:
5am ▾

Morning Activity

Moderate	Hours:Minutes
	0 ▾ :00 ▾
Hard	Hours:Minutes
	0 ▾ :00 ▾
Very Hard	Hours:Minutes
	0 ▾ :00 ▾

Afternoon Activity

Moderate	Hours:Minutes
	0 ▾ :00 ▾
Hard	Hours:Minutes
	0 ▾ :00 ▾
Very Hard	Hours:Minutes
	0 ▾ :00 ▾

Evening Activity

Moderate	Hours:Minutes
	0 ▾ :00 ▾
Hard	Hours:Minutes
	0 ▾ :00 ▾
Very Hard	Hours:Minutes
	0 ▾ :00 ▾

Enter your pedometer reading for the date above:

Activity Log Instructions

Activity Log Help

Please use the guidelines below to help you accurately record your data.

Sleep

This means the time you went to bed one night and the time that you got out of bed the next morning. You may not necessarily have been asleep the entire time you were in bed. You may have been reading or watching television.

For example, if you are filling out the activity log for Monday, you should choose what time you went to sleep Sunday night and what time you woke up Monday morning.

Time of Day

Morning

This is considered from the time you get up in the morning to the time you have lunch.

Afternoon

This includes the time between lunch and dinner.

Night

This is from dinner until the time you go to bed.

Intensity

Moderate

Physical activity that falls into this category is similar to how you feel when you're walking at a brisk pace, walking as if "you're going somewhere." This does *not* include light activities such as desk work, standing, light housework, bowling, strolling, and stop-and-go walking such as grocery or window shopping.

Very Hard

This category is similar to how you feel when you are running.

Hard

This category just falls in between. If the activity seems harder than walking, but not as strenuous as running, it should go in the hard category.

Duration

The activity in question should be performed for a total of at least 10 minutes, intermittently or continuously, during one segment of the day, morning, afternoon, or evening, in which the total amount of minutes is recorded. Make sure that you exclude the time that you stood still or took breaks.

Self-Efficacy Web Form

This questionnaire looks at how confident you are to exercise when other things get in the way. Indicate how confident you would be to exercise in each of the following situations (in your leisure time).

1. When I am under a lot of stress this is how confident I would be to exercise:
2. The roads or sidewalks are snowy
3. I feel I don't have the time
4. I am busy
5. I don't have access to exercise equipment
6. It's cold outside
7. I am alone
8. My exercise partner decides not to exercise that day
9. I don't feel like it
10. I am travelling
11. My gym is closed
12. I am anxious
13. My significant other does not want me to exercise
14. I am feeling down
15. I have to exercise alone
16. I am spending time with friends or family who do not exercise
17. It's raining or snowing
18. My friends don't want me to exercise

Decisional Balance Web Form

This questionnaire looks at positive and negative aspects of exercise. Read the following items and indicate how important each statement is with respect to your decision to exercise or not to exercise in your leisure time. If you disagree with a statement and are unsure how to answer, the statement is probably not important to you.

1. The statement **I would have more energy for my family and friends if I exercised regularly** is this important in my decision whether to exercise or not:
 2. **I would feel less stressed if I exercised regularly**
 3. **I would feel more comfortable with my body if exercised regularly**
 4. **Exercise prevents me from spending time with my friends**
 5. **I would feel embarrassed if people saw me exercising**
 6. **Regular exercise would help me have a more positive outlook on life**
 7. **Exercising puts me in a better mood for the rest of the day**
 8. **I feel uncomfortable or embarrassed in exercise clothes**
 9. **Exercise puts an extra burden on my significant other**
 10. **There is too much I would have to learn to exercise**
-

Working Alliance Inventory Web Forms

■ Please tell us how you feel about Laura (page 1 of 3)

Question	Rating
1. I feel uncomfortable with Laura	<input type="text"/>
2. Laura and I agree about the things I will need to do to help improve my level of physical activity.	<input type="text"/>
3. I am worried about the outcome of my sessions with Laura.	<input type="text"/>
4. What I am doing in my discussions with Laura gives me new ways of looking at physical activity.	<input type="text"/>
5. I understand Laura, and I think she understands me, at least in the best way she can.	<input type="text"/>
6. Laura perceives accurately what my goals are.	<input type="text"/>
7. I find what I am doing with Laura confusing.	<input type="text"/>
8. I feel that Laura, in her own unique way, likes me.	<input type="text"/>
9. I wish Laura and I could clarify the purpose of our sessions.	<input type="text"/>
10. I disagree with Laura about what I ought to get out of my discussions with her.	<input type="text"/>
11. I believe the time Laura and I are spending together is not spent efficiently.	<input type="text"/>
12. Laura does not understand what I am trying to accomplish.	<input type="text"/>

■ Please tell us how you feel about Laura (page 2 of 3)

Question	Rating
13. I am clear on what my responsibilities are with respect to physical activity.	<input type="text"/>
14. My physical activity goals are important to me.	<input type="text"/>
15. I find what Laura and I are doing are unrelated to my concerns.	<input type="text"/>
16. I feel that the things I do with Laura will help me to accomplish the changes that I want.	<input type="text"/>
17. I feel that Laura, in her own unique way, is genuinely concerned about my welfare.	<input type="text"/>
18. I am clear as to what Laura wants me to do in our discussions.	<input type="text"/>
19. I respect Laura, and I feel that she respects me, in her own unique way.	<input type="text"/>
20. I feel that Laura is not totally honest about her feelings toward me.	<input type="text"/>
21. I am confident in Laura's ability to help me.	<input type="text"/>
22. Laura and I are working towards mutually agreed upon goals.	<input type="text"/>
23. I feel that Laura appreciates me.	<input type="text"/>
24. We agree on what is important for me to work on.	<input type="text"/>

■ Please tell us how you feel about Laura (page 3 of 3)

Question	Rating
25. As a result of my discussions with Laura I am clearer as to how I might be able to change.	<input type="text"/>
26. Laura and I trust one another.	<input type="text"/>
27. Laura and I have different ideas on what my problems are.	<input type="text"/>
28. My relationship with Laura is very important to me.	<input type="text"/>
29. I have the feeling that if I say or do the wrong things, Laura will stop working with me.	<input type="text"/>
30. Laura and I collaborate on setting goals for us to work on.	<input type="text"/>
31. I am frustrated by the things I do with Laura.	<input type="text"/>
32. We have established a good understanding of the kind of changes that would be good for me.	<input type="text"/>
33. The things that Laura is asking me to do don't make sense.	<input type="text"/>
34. I don't know what to expect as the result of my discussions with Laura.	<input type="text"/>
35. I believe the way we are working with my problem is correct.	<input type="text"/>
36. I feel Laura cares about me in her own unique way, even when I do things that she does not approve of.	<input type="text"/>

Non-Control Evaluation 1 Web Form

■ Please answer the following questions about the FitTrack system.

Question	Rating
1. How useful were the health tip pages provided to you every day?	<input type="text"/>
2. How useful were the charts of your exercise behavior?	<input type="text"/>
3. How useful were your discussions with Laura?	<input type="text"/>
4. To what extent would you like to continue working with the FitTrack system?	<input type="text"/>
5. To what extent would you like to continue working with Laura?	<input type="text"/>
6. How much do you like Laura?	<input type="text"/>
7. How would you characterize your relationship with Laura?	<input type="text"/>
8. Who was the most helpful to you in getting you to exercise over the last month?	<input type="text"/>
9. Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week. Are you physically active according to that definition?	<input type="text"/>
10. Any other comments or feedback on the FitTrack system or Laura?	<input type="text"/>

Control Evaluation 1 Web Form

■ Please answer the following questions about the FitTrack system.

Question	Rating
1. How useful were the health tip pages provided to you every day?	<input type="text"/>
2. How useful were the charts of your exercise behavior?	<input type="text"/>
3. To what extent would you like to continue working with the FitTrack system?	<input type="text"/>
4. Who was the most helpful to you in getting you to exercise over the last month?	<input type="text"/>
5. Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week. Are you physically active according to that definition?	<input type="text"/>
6. Any other comments or feedback on the FitTrack system?	
<input type="text"/>	

Non-Control Evaluation 2 Web Form

■ Please answer the following questions about the FitTrack system.

Question	Rating
1. To what extent would you like to continue working with the FitTrack system?	<input type="text"/>
2. To what extent would you like to continue working with Laura?	<input type="text"/>
3. Who was the most helpful to you in getting you to exercise over the last month?	<input type="text"/>
4. Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week. Are you physically active according to that definition?	<input type="text"/>
5. Any other comments or feedback on the FitTrack system or Laura?	
<input type="text"/>	

Control Evaluation 2 Web Form

■ Please answer the following questions about the FitTrack system.

Question	Rating
1. To what extent would you like to continue working with the FitTrack system?	<input type="text"/>
2. Who was the most helpful to you in getting you to exercise over the last month?	<input type="text"/>
3. Regular physical activity is any physical activity from walking, slow biking, gardening, or other moderate activities that make you sweat a little, up to and including extremely vigorous exercise, performed a total of 30 minutes or more on most days of the week. Are you physically active according to that definition?	<input type="text"/>
4. Any other comments or feedback on the FitTrack system?	<input type="text"/>

Appendix C: Fragment of Exercise Advisor Scripts

The full set of Exercise Advisor scripts is available as technical report TR 558, from http://vismod.www.media.mit.edu/cgi-bin/tr_pagemaker.

Script Primitives

Type	Function	Purpose
Property Access Functions	GET(<property>)	Looks up single-valued property value in database, returned as a String.
	GETINT(<property>)	Returns value as an integer.
	UNDEFINED(<value>)	Returns 'true' if the property value is undefined.
	SET(<property>,<value>)	Sets a single-valued property in the database.
	GET_SESSION(<property>,<index>)	Looks up historical property value. If (integer) index is positive, this returns the value for the specified session. If index is negative, this uses relative indexing (e.g., "-1" returns the value from the previous session). If zero, this returns the value from the current session.
	GET_SESSION(<property>)	Looks up historical property value for the current session.
	SET_SESSION(<property>,<value>)	Sets a historical property value, indexed to the current session.
	EXISTS_SESSION(<property>,<value>)	Returns most recent session index in which <property> had <value>.
	COUNT_SESSION(<property>,<value>)	Returns the number of sessions in which <property> had value <value>.
	GET_EQ(<property>,<value>)	Returns 'true' if single-valued <property> had value <value>.
	GET_SESSION_EQ(<property>,<value>)	Returns 'true' if historical <property> has value <value> in current session.
	GET_SESSION_EQ(<property>,<index>,<value>)	Returns 'true' if historical <property> has value <value> in the specified session.
ATN State Change Functions	GO(<state>)	Branches to the named state.
	PUSH(<nextState>,<returnState>)	Pushes <returnState> on the runtime stack and branches to <nextState>.
	POP()	Pops the runtime stack and branches to the returned state.
User Interface Functions	GETTEXT()	Returns the text entered by the user.
	CONTENT(<n>)	Displays educational content page for session <n> in the main browser window.
Database lookup	ISREL()	Returns 'true' if user is in RELATIONAL condition of the study, else 'false' (if in NON-RELATIONAL).
	NAME()	Returns subject's given name.

Script Primitives, continued

Type	Function	Purpose
Miscellaneous	PICK(<n>)	Returns a random number between 1 and <n>.
	REFER_SESSION(<n>)	Returns an English expression referring to the time of the specified session (e.g., "YESTERDAY", "LAST TUESDAY", etc.)
	GOALTIME()	Returns the number of minutes the user should be exercising today.
	GOALSTEPS()	Returns the number of steps the user should be walking today.
	ALERT(<string>)	Enters message into the Alert table in the database (reviewed daily by operator for exceptions).

"Int1" script

```

STATE: Int1
AGENT_REL: $ <BEAT> <HAPPY/> Hi there. To talk to me, <TEXTBOX> just click on one
of the options in the menu below </TEXTBOX>. <NEUTRAL/> </BEAT> $
AGENT_NREL: $ <BEAT> <NEUTRAL/> Hi there. To talk to me, <TEXTBOX> just click on
one of the options in the menu below </TEXTBOX>. </BEAT> $
REPEAT: $ <BEAT>Just click the mouse below where it says OK.</BEAT> $
USERMENU: OK, I understand that this is where I click. => $ GO("INT1_1"); $

STATE: INT1_1
ACTION: $ if(ISREL()) GO("INT1_3"); else GO("INT1_2"); $

STATE: INT1_2
AGENT: $ <BEAT> Good. I'm going to be your exercise advisor. </BEAT> $
ACTION: $ GO("INT1_11"); $

STATE: INT1_3
AGENT: $ <BEAT> Good. My name is Laura. <YOU> You are =| GET("NAME") |= , is that right?
</YOU> </BEAT> $
USERMENU: Yep, that's right. => $ GO("INT1_5A"); $
No, actually that's not right. => $ GO("INT1_4"); $

STATE: INT1_4
AGENT: $ <BEAT> Oh, I'm sorry. What is your name? </BEAT> $
USERTEXT: My name is: => $ String entry=GETTEXT();
ALERT("WRONG NAME: not "+GET("NAME")+ " corrected to "+entry);
entry=FILTERNAME(entry);
if(entry.equals("NOTCHECKED")) entry="";
SET("NAME",entry);
GO("INT1_5"); $

STATE: INT1_5A
AGENT: $ <BEAT> Hi, =| GET("NAME") |= . </BEAT> $
ACTION: $ GO("INT1_5"); $

STATE: INT1_5
AGENT: $ <BEAT> It's great to meet you. Sorry about my voice. This is some engineer's idea
of natural sounding. </BEAT> $
USERMENU: It's OK. ; It does sound kind of funny. => $ GO("INT1_6"); $
Yeah. So, can we get down to business? => $ GO("INT1_9"); $

STATE: INT1_6
AGENT: $ <BEAT> Yeah, I know. This whole set up can feel a little strange.
I hope you are more or less comfortable. </BEAT> $
USERMENU: I'm fine, thanks. => $ GO("INT1_8"); $
It's a little strange. => $ GO("INT1_7"); $

STATE: INT1_7

```

AGENT: \$ <BEAT> That's okay. Perhaps it will get more comfortable after we have a few conversations. <NEUTRAL/> </BEAT> \$
ACTION: \$ GO("INT1_9"); \$

STATE: INT1_8
AGENT: \$ <BEAT> <HAPPY/> That's great, I'm glad. <NEUTRAL/></BEAT> \$
ACTION: \$ GO("INT1_9"); \$

STATE: INT1_9
AGENT: \$ <BEAT> So, I'm going to help you meet your exercise goals over the next month. </BEAT> \$
ACTION: \$ GO("INT1_10"); \$

STATE: INT1_10
AGENT: \$ <BEAT> I'm going to be your exercise advisor. My job is to help you set goals and overcome obstacles, and I'll also be giving you lots of tips on how to exercise. </BEAT> \$
ACTION: \$ GO("INT1_11"); \$

STATE: INT1_11
AGENT: \$ <BEAT> You need to keep in mind that I'm just a computer character with limited capabilities, so I hope you will understand if there are some things I cannot help you with. </BEAT> \$
ACTION: \$ if(ISREL()) GO("INT1_12"); else PUSH("GetGoals", "INT1_15"); \$

STATE: INT1_12
AGENT: \$ <BEAT> So, how do you feel about talking to a computer character? </BEAT> \$
USERMENU: It's kind of neat.; It's okay. => \$ GO("INT1_14"); \$
It's kind of strange. => \$ GO("INT1_13"); \$

STATE: INT1_13
AGENT: \$ <BEAT>You are not the only one that feels that way. I hope this gets more comfortable for us as we continue our work together. <NEUTRAL/></BEAT> \$
ACTION: \$ PUSH("GetGoals", "INT1_15"); \$

STATE: INT1_14
AGENT: \$ <BEAT> <HAPPY/>That's good to hear. </BEAT> \$
ACTION: \$ PUSH("GetGoals", "INT1_15"); \$

STATE: INT1_15
AGENT: \$ <BEAT> So, =|" "+(ISREL()?GET("NAME"):"")|= , I'd like you to start thinking about getting out and walking for exercise. <INT_PAUSE DUR="200"/> Before we chat again, I'd like you to just go for a walk. <INT_PAUSE DUR="200"/>It doesn't matter how long or how far. Can you do that for me? </BEAT> \$
REPEAT: \$ <BEAT> Do you think you can go for a short walk before we talk again?</BEAT> \$
USERMENU: Sure, no problem.; I'll try. => \$ GO("INT1_16"); \$
I don't think so. => \$ PUSH("MotivateToWalk", "INT1_16"); \$

STATE: INT1_16
AGENT: \$ =|CONTENT(1)|= <BEAT>OK. <PAGE> Here is some information about walking for exercise. </PAGE> </BEAT> \$
ACTION: \$ GO("INT1_17"); \$

STATE: INT1_17
AGENT: \$ <BEAT> When we end our sessions, you can spend as much time as you like browsing <PAGE>this library</PAGE>. There's lots of good stuff on walking and exercise in general. </BEAT> \$
ACTION: \$ GO("INT1_18"); \$

STATE: INT1_18
AGENT: \$ <BEAT>
If there is something you would like to know that is not in the library, feel free to contact the Health Education department at the MIT medical center. </BEAT> \$
ACTION: \$ GO("INT1_19"); \$

```
STATE: INT1_19
AGENT: $ <BEAT> Just close the FitTrack application when you're done browsing. </BEAT> $
ACTION: $ PUSH("Farewell", "INT1_END"); $
```

```
STATE: INT1_END
ACTION: $ POP(); $
```

"Int2" script

```
STATE: Int2
ACTION: $ PUSH("Greeting", "INT2_17"); $
```

```
STATE: INT2_17
ACTION: $ if(ISREL()) { PUSH("FindState", "INT2_1"); } else { GO("INT2_1"); } $
```

```
STATE: INT2_1
ACTION: $ if(ISREL()) GO("INT2_20"); else GO("INT2_8"); $
```

```
STATE: INT2_20
AGENT: $ <BEAT> So, are you a student? </BEAT> $
USERMENU: Yes => $ SET("OCCUPATION", "STUDENT"); GO("INT2_18"); $
          No => $ GO("INT2_5"); $
```

```
STATE: INT2_18
AGENT: $ <BEAT> Do you go to MIT? </BEAT> $
USERMENU: Yes => $ SET("MIT_STUDENT", "TRUE"); SET("MIT_COMMUNITY", "TRUE"); GO("INT2_2"); $
          No => $ SET("MIT_STUDENT", "FALSE"); SET("MIT_COMMUNITY", "FALSE"); GO("INT2_19"); $
```

```
STATE: INT2_19
AGENT: $ <BEAT> Oh, what school do you go to? </BEAT> $
USERTEXT: I go to: => $ SET("SCHOOL", GETTEXT()); GO("INT2_2"); $
```

```
STATE: INT2_2
AGENT: $ <BEAT> Really? Are you a grad student, or an undergraduate? </BEAT> $
USERMENU:
  Graduate => $ SET("STUDENT_YEAR", "GRADUATE"); GO("INT2_3"); $
  Freshman => $ SET("STUDENT_YEAR", "FRESHMAN"); GO("INT2_3"); $
  Sophomore => $ SET("STUDENT_YEAR", "SOPHOMORE"); GO("INT2_3"); $
  Junior => $ SET("STUDENT_YEAR", "JUNIOR"); GO("INT2_3"); $
  Senior => $ SET("STUDENT_YEAR", "SENIOR"); GO("INT2_3"); $
```

```
STATE: INT2_3
AGENT: $ <BEAT> Cool. What's your major? </BEAT> $
USERTEXT: Enter your major => $ SET("STUDENT_MAJOR", GETTEXT()); GO("INT2_8"); $
```

```
STATE: INT2_5
AGENT: $ <BEAT> What do you do? </BEAT> $
USERMENU:
  I work at MIT. => $ SET("OCCUPATION", "STAFF"); SET("MIT_COMMUNITY", "TRUE"); GO("INT2_6"); $
  I work at another university. => $ SET("OCCUPATION", "STAFF"); GO("INT2_6"); $
  I work in private industry => $ SET("OCCUPATION", "INDUSTRY"); GO("INT2_5A"); $
  I work for the government => $ SET("OCCUPATION", "GOVERNMENT"); GO("INT2_5B"); $
  I'm in between jobs right now => $ SET("OCCUPATION", "UNEMPLOYED"); GO("INT2_5C"); $
  I'm retired => $ SET("OCCUPATION", "RETIRED"); GO("INT2_5D"); $
  I'm a full-time parent => $ SET("OCCUPATION", "PARENT"); GO("INT2_5E"); $
```

```
STATE: INT2_6
AGENT: $ <BEAT> Really? What department? </BEAT> $
USERTEXT: My department is: => $ SET("MIT_DEPARTMENT", GETTEXT()); GO("INT2_7"); $
```

```
STATE: INT2_5A
AGENT: $ <BEAT> Private industry? That's great. </BEAT> $
ACTION: $ GO("INT2_7"); $
```

```
STATE: INT2_5B
AGENT: $ <BEAT> So you work in the public sector. That's great. </BEAT> $
ACTION: $ GO("INT2_7"); $
```

STATE: INT2_5C
AGENT: \$ <BEAT> When you're in between jobs, sometimes you have a lot of time to take care of yourself and do things like exercise. So that's great. </BEAT> \$
ACTION: \$ GO("INT2_7"); \$

STATE: INT2_5D
AGENT: \$ <BEAT>So you're retired. You probably have a flexible schedule, which is great for a fitness program. </BEAT> \$
ACTION: \$ GO("INT2_7"); \$

STATE: INT2_5E
AGENT: \$ <BEAT> That's a <EYEBROW DIR="UP"/>tough job! <EYEBROW DIR="DOWN"/> Maybe you can get the kids involved in your exercise program. </BEAT> \$
ACTION: \$ GO("INT2_7"); \$

STATE: INT2_7
AGENT: \$ <BEAT>Well, I work as an exercise advisor at many locations around Boston.</BEAT> \$
ACTION: \$ GO("INT2_8"); \$

STATE: INT2_8
AGENT: \$ <BEAT> So, now I'd like to ask you a little about your exercise behavior.
Do you currently do any kind of sport? </BEAT> \$
USERMENU:
No, I don't => \$ GO("INT2_10"); \$
I play squash => \$ SET("SPORTPLAYED", "SQUASH"); GO("INT2_10"); \$
I play golf => \$ SET("SPORTPLAYED", "GOLF"); GO("INT2_10"); \$
I play tennis => \$ SET("SPORTPLAYED", "TENNIS"); GO("INT2_10"); \$
I row => \$ SET("SPORTPLAYED", "ROW"); GO("INT2_10"); \$
I play hockey => \$ SET("SPORTPLAYED", "HOCKEY"); GO("INT2_10"); \$
I play softball => \$ SET("SPORTPLAYED", "SOFTBALL"); GO("INT2_10"); \$
I play basketball => \$ SET("SPORTPLAYED", "BASKETBALL"); GO("INT2_10"); \$
I play soccer => \$ SET("SPORTPLAYED", "SOCCER"); GO("INT2_10"); \$
I play another sport => \$ SET("SPORTPLAYED", "OTHER"); GO("INT2_9"); \$

STATE: INT2_9
AGENT: \$ <BEAT> Oh, what sport do you do? </BEAT> \$
USERTEXT: My sport is: => \$ SET("SPORT_OTHER", GETTEXT()); GO("INT2_10"); \$

STATE: INT2_10
AGENT: \$ <BEAT> Do you currently do any other kind of aerobic exercise regularly? </BEAT> \$
USERMENU:
I go to the gym => \$ SET("AEROBIC_EXERCISE", "GYM"); GO("INT2_12"); \$
I like to run => \$ SET("AEROBIC_EXERCISE", "RUN"); GO("INT2_12"); \$
I go for walks => \$ SET("AEROBIC_EXERCISE", "WALK"); GO("INT2_12"); \$
I swim => \$ SET("AEROBIC_EXERCISE", "SWIM"); GO("INT2_12"); \$
I bike => \$ SET("AEROBIC_EXERCISE", "BIKE"); GO("INT2_12"); \$
I do something else => \$ GO("INT2_11"); \$
Nope => \$ GO("INT2_12"); \$

STATE: INT2_11
AGENT: \$ <BEAT> What do you do? </BEAT> \$
USERTEXT: For aerobic exercise, I: => \$ SET("AEROBIC_OTHER", GETTEXT()); GO("INT2_12"); \$

STATE: INT2_12
AGENT: \$ <BEAT> And, do you
=|(UNDEFINED(GET("AEROBIC_EXERCISE"))?)?" think you would ":" usually ")|=
like to exercise by yourself, or with a buddy? </BEAT> \$
USERMENU: By myself => \$ SET("ISLONER", "TRUE"); GO("INT2_13"); \$
With a buddy => \$ SET("ISLONER", "FALSE"); GO("INT2_13"); \$

STATE: INT2_13
AGENT: \$ <BEAT> Great. So, did you manage to go for a walk since we last talked?</BEAT> \$
USERMENU: Yes! => \$ GO("INT2_13B"); \$
No, didn't get the chance. => \$ GO("INT2_13C"); \$

STATE: INT2_13B
AGENT: \$ <BEAT> <HAPPY/><OK>That's wonderful.<INT_PAUSE DUR="500"/></OK><NEUTRAL/></BEAT> \$

```

ACTION: $ GO("INT2_15"); $

STATE: INT2_13C
AGENT: $ <BEAT> OK. </BEAT> $
ACTION: $ GO("INT2_15"); $

STATE: INT2_15
AGENT: $ <BEAT> Before we chat again, do you think you could get out and go for another walk?
It doesn't matter how long or how far. </BEAT> $
USERMENU: Sure. => $ GO("INT2_14"); $
          No, I don't think so. => $ PUSH("MotivateToWalk", "INT2_14"); $

STATE: INT2_14
AGENT: $ =|CONTENT(2)|= <BEAT> Okay. <PAGE>Here is some information about how to make a walk
a workout. </PAGE>
          Casual strolls don't quite count. </BEAT> $
ACTION: $ GO("INT2_16"); $

STATE: INT2_16
AGENT: $ <BEAT> Remember, when we end our sessions, you can spend as much time as you like
browsing this library. Just close the Fit Track application when you are finished. </BEAT>
$
ACTION: $ PUSH("Farewell", "INT2_END"); $

STATE: INT2_END
ACTION: $ POP(); $

```

"Int3" script

```

STATE: Int3
ACTION: $ PUSH("Greeting", "INT3_1"); $

STATE: INT3_1
ACTION: $ if(ISREL()) PUSH("FindState", "INT3_2"); else GO("INT3_2"); $

STATE: INT3_2
ACTION: $ if(ISREL()) {
          if(GET_EQ("OCCUPATION", "STUDENT"))
            GO("INT3_14");
          else GO("INT3_4");
        }
        else GO("INT3_20"); $

STATE: INT3_14
AGENT: $ <BEAT> So, <INT_EMPHASIS/>how is your <INT_EMPHASIS/>semester going so far? </BEAT> $
USERMENU: Great; Okay => $ GO("INT3_14A"); $
          Not so good; Pretty bad => $ GO("INT3_14B"); $
          Actually, I'm in a bit of a rush today. => $ INCREMENT_STAT(1); GO("INT3_20"); $

STATE: INT3_14A
AGENT: $ <BEAT> That's good to hear. </BEAT> $
ACTION: $ GO("INT3_3"); $

STATE: INT3_14B
AGENT: $ <BEAT> I'm sorry to hear that. </BEAT> $
ACTION: $ GO("INT3_3"); $

STATE: INT3_3
AGENT: $ <BEAT> Are you <INT_EMPHASIS/>taking a full <INT_EMPHASIS/>course load? </BEAT> $
USERMENU: Yes => $ GO("INT3_6"); $
          Nope, I'm going part time. => $ GO("INT3_6"); $

STATE: INT3_6
AGENT: $ <BEAT> Do you enjoy
          =|""+(GET_EQ("MIT_STUDENT", "TRUE"))? " MIT?" : " your school?"| = </BEAT> $
USERMENU: It's great. How about you? => $ GO("INT3_7"); $

```

```

        I don't like it much. => $ GO("INT3_7"); $

STATE: INT3_7
AGENT: $ <BEAT> I can't wait to check out the new athletic center. Of course, I can't
        personally use it because I don't have any legs. </BEAT> $
ACTION: $ GO("INT3_20"); $

STATE: INT3_4
AGENT: $ <BEAT> So, how do you like Boston? </BEAT> $
USERMENU: I like it a lot. => $ GO("INT3_4B"); $
        I don't like it much. => $ GO("INT3_4C"); $
        Actually, I'm in a bit of a rush today. => $ INCREMENT_STAT(1); GO("INT3_20"); $

STATE: INT3_4B
AGENT: $ <BEAT> Really? That's great. </BEAT> $
ACTION: $ GO("INT3_5"); $

STATE: INT3_4C
AGENT: $ <BEAT> Really? That's too bad. </BEAT> $
ACTION: $ GO("INT3_5"); $

STATE: INT3_5
AGENT: $ <BEAT> What about New <INT_EMPHASIS/>England? </BEAT> $
USERMENU: It's great. => $ GO("INT3_5B"); $
        It's not that great. => $ GO("INT3_5C"); $

STATE: INT3_5B
AGENT: $ <BEAT> Cool. Boston is the only home I've ever known. </BEAT> $
ACTION: $ if(GET_EQ("MIT_COMMUNITY","TRUE")) GO("INT3_6"); else GO("INT3_20"); $

STATE: INT3_5C
AGENT: $ <BEAT> Sorry to hear that. Boston is the only home I've ever known. </BEAT> $
ACTION: $ if(GET_EQ("MIT_COMMUNITY","TRUE")) GO("INT3_6"); else GO("INT3_20"); $

STATE: INT3_20
AGENT: $ <DELAY MS="500"/> <BEAT> So, did you get that walk in?</BEAT> $
USERMENU: Yes! => $ GO("INT3_21"); $
        No, didn't get the chance. => $ GO("INT3_22"); $

STATE: INT3_21
AGENT: $ <BEAT> <HAPPY/><OK>Way to go!<INT_PAUSE DUR="500"/></OK><NEUTRAL/></BEAT> $
ACTION: $ GO("INT3_10"); $

STATE: INT3_22
AGENT: $ <BEAT> OK. </BEAT> $
ACTION: $ GO("INT3_10"); $

STATE: INT3_10
ACTION: $ if (ISREL()) GO("INT3_11"); else GO("INT3_8"); $

STATE: INT3_11
AGENT: $ <BEAT> So, how are you feeling about working with me? Are you comfortable with this?
</BEAT> $
USERMENU: Great.; No problem.; => $ GO("INT3_13"); $
        Actually, it's a little strange.; Well, it's not the best so far.; => $
GO("INT3_12"); $
        I'm afraid of letting you down. => $ GO("INT3_15"); $

STATE: INT3_15
AGENT: $ <BEAT> The most important thing to me, is that you just log-in every day, and just
keep trying.
        And I know you will do great. </BEAT> $
ACTION: $ GO("INT3_8"); $

STATE: INT3_12
AGENT: $ <BEAT> I am sorry to hear that. I do want to help you as much as I can. </BEAT> $
ACTION: $ GO("INT3_8"); $

STATE: INT3_13

```

```

AGENT: $ <BEAT> That's great. I really enjoy the time we spend together, too. </BEAT> $
ACTION: $ GO("INT3_8"); $

STATE: INT3_8
AGENT: $ <BEAT> I'm looking forward to helping you reach your exercise goals. If you get the
      chance to get out and go for another walk before we talk again that would be
great. </BEAT> $
ACTION: $ GO("INT3_9"); $

STATE: INT3_9
AGENT: $ =|CONTENT(3)|= <BEAT> So, today I'd like you to <PAGE> read this note about
recommendations for
      minimum physical activity </PAGE>, because the next time we chat I'd like you to set a goal
for the next month, OK? </BEAT> $
ACTION: $ PUSH("Farewell", "INT3_END"); $

STATE: INT3_END
ACTION: $ POP(); $

```

"Int4" script

```

STATE: Int4
ACTION: $ PUSH("Greeting", "INT4_1"); $

STATE: INT4_1
ACTION: $ if (ISREL()) PUSH("FindState", "INT4_3"); else GO("INT4_3"); $

STATE: INT4_3
ACTION: $
  if (!ISREL())
    GO("INT4_14");
  else {
    if (GET_EQ("OCCUPATION", "STUDENT"))
      GO("INT4_5");
    else
      GO("INT4_4");
  } $

STATE: INT4_4
AGENT: $ <BEAT>So, where do you live? </BEAT> $
USERMENU:
In Cambridge. => $ SET("RESIDENCE", "CAMBRIDGE"); GO("INT4_8B"); $
In Somerville => $ SET("RESIDENCE", "SOMERVILLE"); GO("INT4_8B"); $
In Boston => $ SET("RESIDENCE", "BOSTON"); GO("INT4_8B"); $
You'll never guess. => $ GO("INT4_7"); $
Uh, whatever. => $ INCREMENT_STAT(1); GO("INT4_14");$

STATE: INT4_7
AGENT: $ <BEAT> Where is that? </BEAT> $
REPEAT: $ <BEAT> Where do you live? </BEAT> $
USERTEXT: I live in: => $ SET("RESIDENCE_OTHER", GETTEXT()); GO("INT4_8B"); $

STATE: INT4_5
AGENT: $ <BEAT> So, where do you live? </BEAT> $
REPEAT: $ <BEAT> where do you live? </BEAT> $
USERMENU:
  In a dorm. => $ SET("STUDENT_RESIDENCE", "DORM"); GO("INT4_8A"); $
  In a fraternity. => $ SET("STUDENT_RESIDENCE", "FRATERNITY"); GO("INT4_8A"); $
  In a sorority. => $ SET("STUDENT_RESIDENCE", "SORORITY"); GO("INT4_8A"); $
  In a living group. => $ SET("STUDENT_RESIDENCE", "LIVINGGROUP"); GO("INT4_8A"); $
  Off campus. => $ SET("STUDENT_RESIDENCE", "OFFCAMPUS"); GO("INT4_6"); $
  Uh, whatever. => $ GO("INT4_14"); $

STATE: INT4_6
AGENT: $ <BEAT> <INT_EMPHSAIS/>Where off campus?</BEAT> $
REPEAT: $ <BEAT> <INT_EMPHSAIS/>Where off campus do you live?</BEAT> $

```

```

USERTEXT: Where I live is: => $ GO("INT4_8A"); $

STATE: INT4_8A // students
AGENT: $
<BEAT>=|(""+(UNDEFINED(GET("STUDENT_RESIDENCE"))?"":LOC_PP(GET("STUDENT_RESIDENCE"))))|=?
    Really? </BEAT> $
USERMENU: Yeah. And you? => $ GO("INT4_9"); $
    Yeah, it's great.; Yeah, it's OK.; Yeah, it's not that great. => $ GO("INT4_10"); $

STATE: INT4_8B // non students
AGENT: $ <BEAT>=|(""+(UNDEFINED(GET("RESIDENCE"))?"":LOC_PP(GET("RESIDENCE"))))|=?
    Really? </BEAT> $
USERMENU: Yeah. And you? => $ GO("INT4_9"); $
    Yeah, it's great.; Yeah, it's OK.; Yeah, it's not that great. => $ GO("INT4_10"); $

STATE: INT4_9
AGENT: $ <BEAT>Well it's hard to say. You see I live in many places at once.
    It's kind of interesting being able to be replicated. </BEAT> $
ACTION: $ GO("INT4_10"); $

STATE: INT4_10
AGENT: $ <BEAT>So, do you live by yourself? </BEAT> $
REPEAT: $ <BEAT>do you live alone? </BEAT> $
USERMENU: Yes => $ GO("INT4_12");$
    No, I have a roommate.;
    No, I have roommates.;
    No, I live with my partner;
    No, I live with my family
=> $ GO("INT4_13");$

STATE: INT4_12
AGENT: $ <BEAT>All to yourself. Cool. </BEAT> $
ACTION: $ GO("INT4_14");$

STATE: INT4_13
AGENT: $ <BEAT>That's great. Any potential exercise buddies? </BEAT> $
REPEAT: $ <BEAT>Any potential exercise buddies at home? </BEAT> $
USERMENU: Yeah, sure; Maybe => $ GO("INT4_13A"); $
    I don't think so => $ GO("INT4_13B"); $

STATE: INT4_13A
AGENT: $ <BEAT>Cool. Having a buddy to exercise with can make it a lot of fun. </BEAT> $
ACTION: $ GO("INT4_14"); $

STATE: INT4_13B
AGENT: $ <BEAT> That's okay. If you like exercising with a buddy, I'm sure you can
    find one somewhere else. </BEAT> $
ACTION: $ GO("INT4_14"); $

STATE: INT4_14
AGENT_REL: $ <BEAT>So, let's set a goal for you to work towards this month. </BEAT> $
AGENT_NREL: $ <BEAT>So, I'd like you to set a goal to work towards this month. </BEAT> $
REPEAT: $ <BEAT>It's time to set a goal for you to work towards this month. </BEAT> $
USERMENU: OK. => $ GO("INT4_16");$
    I'd rather not. => $ GO("INT4_15"); $

STATE: INT4_15
AGENT_REL: $ <BEAT> It's <INT_EMPHASIS/>really important to have something to
<INT_EMPHASIS/>work towards. Can you help me
    out on this?</BEAT> $
AGENT_NREL: $ <BEAT> It's <INT_EMPHASIS/>really important to have something to
<INT_EMPHASIS/>work towards. Now is the right time in your program to set a goal. </BEAT> $
USERMENU: If you insist. => $ GO("INT4_16");$

STATE: INT4_16
AGENT: $ <BEAT>You've been averaging
    =| GET("GOAL_TIME") |=
    minutes a day of moderate or better physical activity.

```

```

        The guidelines that you red last time suggest that you should be getting 30 minutes
        a day. </BEAT> $
ACTION:
$
    if (GETINT("GOAL_TIME") >= 30)
        GO("INT4_17");
    else
        GO("INT4_18");
$

STATE: INT4_17
AGENT: $ <BEAT>Since you're already above the goal level, why don't we just try to maintain
        your current level? </BEAT> $
ACTION: $ PUSH("GetCommitment", "INT4_22");$

STATE: INT4_18
AGENT: $ <BEAT>Why don't we try to work towards 30 minutes a day of walking, or other
        physical activity, as a goal? </BEAT> $
ACTION: $ GO("INT4_19");$

STATE: INT4_19
AGENT: $ <BEAT>This isn't something you are going to reach right away, but gradually
        work up to. </BEAT> $
USERMENU: Okay, sounds good. => $ GO("INT4_21"); $
        I still think that sounds like too much. => $ GO("INT4_20");$

STATE: INT4_20
AGENT: $ <BEAT>OK, then we'll just start out with something very easy and go from there.
</BEAT> $
ACTION: $ PUSH("GetCommitment", "INT4_22"); $

STATE: INT4_21
AGENT: $ <BEAT> Great, so let's try to gradually work up to a goal of 30 minutes a
day.</BEAT>$
ACTION: $PUSH("GetCommitment", "INT4_22"); $

STATE: INT4_22
AGENT: $ =|CONTENT(4)|= <BEAT> <PAGE> Here is some good information about ped ometers that I
hope you will find
        useful. </PAGE> I'd like you to pay particular attention to the health benefits
        of walking 10,000 steps a day. </BEAT>$
ACTION: $ PUSH("Farewell", "INT4_END"); $

STATE: INT4_END
ACTION: $ POP();$

```

"Int30" script

```

//Used for final interaction regardless of actual number of logins.

STATE: Int30
ACTION: $ PUSH("Greeting", "INTN_1"); $

STATE: INTN_1
ACTION: $ if (ISREL()) PUSH("FindState", "INTN_2"); else GO("INTN_3"); $

STATE: INTN_2
ACTION: $ PUSH("PastEventTalk", "INTN_3"); $

STATE: INTN_3
ACTION: $ PUSH("FollowUp", "INTN_4"); $

STATE: INTN_4
AGENT: $ <BEAT>So, this will be our last day together. </BEAT> $
ACTION: $ GO("INTN_5"); $

STATE: INTN_5

```

```

AGENT: $ <BEAT>Don't forget that even without the Fit Track program,
      you can keep track of your exercise. Just keep a piece
      of graph paper on the wall with your daily steps or
      time spent on physical activity. </BEAT> $
ACTION: $ GO("INTN_6"); $

STATE: INTN_6
AGENT: $ <BEAT>And, don't forget that even if you can't meet your goal for a
      few days, you can always get right back in to the hab it, even by
      lowering your goal if necessary. The important thing
      is to just keep exercising. </BEAT> $
ACTION: $ GO("INTN_7"); $

STATE: INTN_7
ACTION: $ GO("INTN_8"); $

STATE: INTN_8
AGENT: $ <BEAT> It has been a real pleasure working with you. I hope you are
      able to keep up with your exercise. </BEAT> $
ACTION: $ GO("INTN_9"); $

STATE: INTN_9
AGENT: $ <BEAT>Good luck, and take care =|(ISREL())? GET("NAME"):")|= . </BEAT> $
USERMENU: Bye. => $ GO("INTN_END"); $
          Take care Laura, I'll miss you. => $ INCREMENT_STAT(2);
          if(ISREL()) GO("INTN_10"); else GO("INTN_END"); $

//REL only:
STATE: INTN_10
AGENT: $ <FACE EXPR="CONCERN"/><CAMERA ZOOM="1.0"/><DELAY MS="1000"/>
      <SPEECH><EYEBROWS DIR="UP"/>Thanks, =|GET("NAME")|= \Pau=500\
      I'll miss you too.</SPEECH><DELAY MS="1000"/> $
ACTION: $ GO("INTN_11"); $

//REL only:
STATE: INTN_11
AGENT: $ <FACE EXPR="WARM"/><SPEECH>Well.</SPEECH><DELAY MS="500"/>
      <FACE EXPR="SMILE"/><SPEECH>We had some fun together.</SPEECH>
      <FACE EXPR="WARM"/><SPEECH>Maybe we'll cross paths again someday. \Pau=500\
      Take care of yourself, =|GET("NAME")|= .</SPEECH>
      <DELAY MS="500"/><CAMERA ZOOM="0.6"/>
      <DELAY MS="500"/><CAMERA ZOOM="0.3"/>
      <DELAY MS="500"/><CAMERA ZOOM="0.0"/> $
ACTION: $ GO("INTN_END"); $

//NOTE: Following this interaction the "YOU ARE FINISHED" page is displayed, so
//      dont' bring up a content page now.

STATE: INTN_END
AGENT: $ =|CONTENT(-3)|= $ //wait for followup page
ACTION: $ POP(); $

```

"Farewell" script

```

STATE: Farewell
ACTION: $ if(ISREL()) GO("FW_2"); else GO("FW_5"); $

STATE: FW_2
AGENT: $ <BEAT> So, when can we get together again? </BEAT> $
      $ <BEAT> So, when can we chat again? </BEAT> $
      $ <BEAT> So, when will I see you again? </BEAT> $
      $ <BEAT> So, when do you want to get together next? </BEAT> $
      $ <BEAT> So, I hope we can see each other soon. When can we chat again? </BEAT> $
USERMENU: Tomorrow. => $ GO("FW_3"); $
          Day after tomorrow.; In three days.; In a few days. => $ GO("FW_4"); $

```

```

STATE: FW_4
AGENT: $ <BEAT> Okay =| GET("NAME") |= , take care. </BEAT> $
      $ <BEAT> All right, =| GET("NAME") |= , see you then. </BEAT> $
      $ <BEAT> okay then, =| GET("NAME") |= , I look forward to it. </BEAT> $
      $ <BEAT> okay then, =| GET("NAME") |= , I'll see you in a littlewhile. </BEAT> $ // tts
      $ <BEAT> okay then, =| GET("NAME") |= , I'll be here. Just come by any time. </BEAT> $
ACTION: $ GO("FW_5"); $

STATE: FW_3
AGENT: $ <BEAT> Great. Take care =| GET("NAME") |= , I'll talk to you tomorrow. </BEAT> $
      $ <BEAT> Okay then, =| GET("NAME") |= , see you tomorrow. </BEAT> $
      $ <BEAT> Wonderful, =| GET("NAME") |=. <INT_EMPHASIS/>See you <INT_EMPHASIS/>
          tomorrow then! </BEAT> $
      $ <BEAT> Great. Take care, =| GET("NAME") |= , I'll talk to you tomorrow. </BEAT> $
      $ <BEAT> <INT_EMPHASIS/> Tomorrow it <INT_EMPHASIS/> is <INT_PAUSE DUR="100"/>then,
      =| GET("NAME") |= ! Take care. </BEAT> $
REPEAT: $ <BEAT> <INT_EMPHASIS/><INT_EMPHASIS/>See you tomorrow, =| GET("NAME") |= ! </BEAT> $
ACTION: $ GO("FW_5"); $

STATE: FW_5
ACTION: $ if(GET_EQ("PART_OF_DAY", "MORNING")) GO("FW_9");
          else if (GET_EQ("PART_OF_DAY", "EVENING")) GO("FW_7");
          else GO("FW_6"); $

STATE: FW_9 //morning
AGENT: $ <BEAT> <HAPPY/><YOU>Have a great morning, and the rest of your day too.
Goodbye.</YOU> </BEAT> $
ACTION: $ GO("FW_END"); $

STATE: FW_6 //afternoon
AGENT: $ <BEAT> <HAPPY/><YOU>Have a great day. Goodbye.</YOU> </BEAT> $
ACTION: $ GO("FW_END"); $

STATE: FW_7 //evening
AGENT: $ <BEAT> <HAPPY/><YOU>Good <INT_EMPHASIS/>night.</YOU> </BEAT> $
      $ <BEAT> <HAPPY/><YOU>Have a good <INT_EMPHASIS/>evening.</YOU> </BEAT> $
ACTION: $ GO("FW_END"); $

STATE: FW_END
ACTION: $ POP(); $

```

"FindState" script

```

//Only called in RELATIONAL

STATE: FindState
AGENT: $ <BEAT>How are you doing today? </BEAT> $
      $ <BEAT>How are you feeling today? </BEAT> $
      $ <BEAT>How is it going? </BEAT> $
USERMENU: Great! How are you?; All right. And you? => $ GO("FS_11"); $
          So-so.; I'm not feeling so great. => $ GO("FS_7"); $

STATE: FS_7
AGENT: $ <BEAT> I'm sorry to hear that. </BEAT> $
ACTION: $
      if(GET_SESSION_EQ("STATE", -1, "SICK"))
        GO("FS_9");
      else if(GET_SESSION_EQ("STATE", -1, "INJURED"))
        GO("FS_10");
      else if(GET_SESSION_EQ("STATE", -1, "UPSET"))
        GO("FS_12");
      else if(GET_SESSION_EQ("STATE", -1, "DOWN"))
        GO("FS_14");
      else if(GET_SESSION_EQ("STATE", -1, "TIRED"))
        GO("FS_15");
      else if(GET_SESSION_EQ("STATE", -1, "STRESSED"))
        GO("FS_16");

```

```

else /* no prior state, or other */
    GO("FS_1");
$

STATE: FS_1
AGENT: $ <BEAT>Tell me more about how you are feeling </BEAT> $
USERMENU: I'm sick => $ SET_SESSION("STATE", "SICK");
    PUSH("NewIllness", "FS_17"); $
    I hurt myself. => $ SET_SESSION("STATE", "INJURED");
        PUSH("NewInjury", "FS_17"); $
    I'm tired. => $ SET_SESSION("STATE", "TIRED");
        PUSH("NewTired", "FS_17"); $
    I'm feeling down. => $ SET_SESSION("STATE", "DOWN");
        PUSH("NewFeelingDown", "FS_17"); $
    I'm feeling upset. => $
    SET_SESSION("STATE", "UPSET");
    PUSH("NewFeeling", "FS_17"); $
    I'm a little STRESSED OUT. => $
    SET_SESSION("STATE", "STRESSED");
    PUSH("NewStressed", "FS_17"); $
    I'm OK.; I don't want to talk about it now. => $ GO("FS_17");$

STATE: FS_9
AGENT: $ <BEAT> Are you still feeling sick? </BEAT> $
REPEAT: $ <BEAT> Are you still feeling poorly? </BEAT> $
USERMENU: Yes, I am. => $ SET_SESSION("STATE", "SICK");
    PUSH("ContinuingIllness", "FS_17"); $
    No, I'm feeling better. => $ GO("FS_13"); $

STATE: FS_10
AGENT: $ <BEAT> Are you still hurt? </BEAT> $
REPEAT: $ <BEAT> Are you still injured? </BEAT> $
USERMENU: Yes, I am. => $ SET_SESSION("STATE", "INJURED");
    PUSH("ContinuingInjury", "FS_17"); $
    No, I'm feeling better. => $ GO("FS_13"); $

STATE: FS_12
AGENT: $ <BEAT> Are you still feeling
    =| "+"(UNDEFINED(GET_SESSION("FEELING", -1))? "upset" : GET_SESSION("FEELING", -1)) |=
    ? </BEAT> $
USERMENU: Yes, I am. =>
    $ SET_SESSION("STATE", "UPSET"); PUSH("ContinuingFeeling", "FS_17"); $
    No, I'm feeling better. => $ GO("FS_13"); $

STATE: FS_13
AGENT: $ <BEAT>That's good. </BEAT> $
ACTION: $ GO("FS_1");$

STATE: FS_11
AGENT: $ <BEAT> Fine thanks. </BEAT> $
    $ <BEAT> Same as usual. </BEAT> $
    $ <BEAT> I'm doing great thanks. </BEAT> $
    $ <BEAT> Oh, just fine thanks. </BEAT> $
ACTION: $ GO("FS_END");$

STATE: FS_14
AGENT: $ <BEAT> Are you still feeling down? </BEAT> $
USERMENU: Yes, I am. => $ SET_SESSION("STATE", "DOWN"); PUSH("ContinuingFeelingDown", "FS_17");
$
    No, I'm feeling better. => $ GO("FS_13"); $

STATE: FS_15
AGENT: $ <BEAT> Are you still tired? </BEAT> $
USERMENU: Yes, I am. => $ SET_SESSION("STATE", "TIRED"); PUSH("ContinuingTired", "FS_17"); $
    No, I got some rest. => $ GO("FS_13"); $

STATE: FS_16
AGENT: $ <BEAT> Are you still feeling stressed out? </BEAT> $
USERMENU: YES! => $ SET_SESSION("STATE", "STRESSED"); PUSH("ContinuingStressed", "FS_17"); $

```

```

        No, I'm a little more relaxed today. => $ GO("FS_13"); $

STATE: FS_17
AGENT: $ <DELAY MS="1500"/><FACE EXPR="WARM"/><GAZE DIR="AWAY"/><DELAY MS="1500"/> $
ACTION: $ GO("FS_END"); $

STATE: FS_END
AGENT: $ <BEAT> <NEUTRAL/> OK. </BEAT> $
ACTION: $ POP(); $

```

"FollowUp" script

```

STATE: FollowUp
AGENT: $ =|CONTENT(-2)|= $
ACTION: $ if(UNDEFINED(GETINT("COMMITMENT_TIME")) && GETINT("EXERCISE_TIME")<10) GO("FU_9");
        else GO("FU_1B"); $

STATE: FU_1B
AGENT: $ <BEAT><NEUTRAL/> So, let's talk about how you did since the last time we got
together.</BEAT> $
        $ <BEAT><NEUTRAL/> So, let's review your progress.</BEAT> $
        $ <BEAT><NEUTRAL/> So, let's talk about how you did on your exercise.</BEAT> $
ACTION: $ GO("FU_1C"); $

STATE: FU_1C
AGENT: $ <BEAT>In order to meet <PAGE> your long-term goal </PAGE>,
        you should be exercising =|GET("LAST_GOAL_TIME")|= minutes a day now. </BEAT>$
REPEAT: $ <BEAT> You should be exercising =|GET("LAST_GOAL_TIME")|= minutes a day now.,
        in order to meet <PAGE> your long-term goal </PAGE>.</BEAT>$
ACTION: $ if(!UNDEFINED(GET("COMMITMENT_TIME"))) GO("FU_1D"); else GO("FU_1E"); $

STATE: FU_1D
AGENT: $ <BEAT>And, you said you would exercise for <PAGE> =|GET("COMMITMENT_TIME")|= minutes
</PAGE>
        the last           time we spoke. </BEAT> $
AGENT: $ <BEAT>And, you told me you would exercise for <PAGE> =|GET("COMMITMENT_TIME")|=
minutes </PAGE>. </BEAT> $
ACTION: $ GO("FU_1E"); $

STATE: FU_1E
ACTION: $   int commitment=GETINT("COMMITMENT_TIME");
           int goal=GETINT("LAST_GOAL_TIME");
           int actual=GETINT("EXERCISE_BEST_TIME");
           boolean madeCommitment=!UNDEFINED(commitment);
           boolean didSomeExercise=(actual>=10);
           if(!didSomeExercise)
               GO("FU_6");
           else if(madeCommitment) {
               if(actual>=commitment)
                   GO("FU_1");
               else if(actual<commitment && actual<goal)
                   GO("FU_2");
               else
                   GO("FU_3");
           } else { /* no commitment made */
               if(actual>=goal)
                   GO("FU_4");
               else
                   GO("FU_5");
           };
           $

STATE: FU_1
AGENT: $ <BEAT><HAPPY/><OK><INT_PITCH_FREQ="220"/>Congratulations.
        <INT_RESET/>Looks like mission accomplished on the exercise.</OK>

```

```

    You met your commitment of =|GET("COMMITMENT_TIME")|= minutes.
    <NEUTRAL/> </BEAT> $
    $ <BEAT><HAPPY/><OK>Looks like you met your exercise commitment of
=|GET("COMMITMENT_TIME")|= minutes,
    <INT_PITCH FREQ="200"/>that's great. <NEUTRAL/></OK></BEAT> $
    $ <BEAT><HAPPY/><OK>Looks like you got your workout in and met your commitment
of =|GET("COMMITMENT_TIME")|= minutes.
    <INT_PITCH FREQ="200"/>Wonderful.</OK> <NEUTRAL/></BEAT> $
ACTION: $ if (ISREL()) GO("FU_1A"); else GO("FU_8"); $

STATE: FU_1A
AGENT: $ <BEAT> We make a great team. </BEAT> $
    $ <BEAT> We're doing some great work together. </BEAT> $
    $ <BEAT> We work together well, I think. </BEAT> $
ACTION: $ GO("FU_8"); $

STATE: FU_2
AGENT: $ <BEAT><PAGE>Looks like you managed to get some exercise in. </PAGE></BEAT> $
    $ <BEAT><PAGE>I see you managed do some exercise. </PAGE></BEAT> $
ACTION: $ GO("FU_8"); $

STATE: FU_3
AGENT: $ <BEAT><PAGE>Looks like you met the exercise goal of =|GET("LAST_GOAL_TIME")|=
minutes. </PAGE></BEAT> $
    $ <BEAT><PAGE>Looks like you got your =|GET("LAST_GOAL_TIME")|= minutes of exercise
in.
    You met your goal.</PAGE></BEAT> $
    $ <BEAT><PAGE>I see you got your =|GET("LAST_GOAL_TIME")|= minutes of exercise in. You
met your goal.</PAGE></BEAT> $
ACTION: $ GO("FU_8");$

STATE: FU_4
AGENT: $ <BEAT><INT_PITCH FREQ="220"/>Congratulations. <INT_RESET/>
    <PAGE>Looks like you met the exercise goal of =|GET("LAST_GOAL_TIME")|= minutes.
</PAGE> </BEAT> $
    $ <BEAT><PAGE>Looks like you got your =|GET("LAST_GOAL_TIME")|= minutes of exercise
in.</PAGE> <INT_PITCH FREQ="200"/> That's <INT_EMPHASIS/> great. </BEAT> $
ACTION: $ GO("FU_7"); $

STATE: FU_5
AGENT: $ <BEAT><PAGE>Looks like you managed to get some exercise in. </PAGE></BEAT> $
    $ <BEAT><PAGE>I see you managed to get some exercise in. </PAGE></BEAT> $
    $ <BEAT><PAGE>I see you managed to do some exercise. </PAGE></BEAT> $
ACTION: $ GO("FU_7"); $

STATE: FU_6
AGENT: $ <BEAT><PAGE>I guess you didn't get a chance to exercise. </PAGE></BEAT> $
    $ <BEAT><PAGE>Looks like you didn't get a chance to exercise. </PAGE></BEAT> $
    $ <BEAT><PAGE>It seems that you didn't get a chance to exercise. </PAGE></BEAT> $
ACTION: $ GO("FU_9"); $

STATE: FU_7
ACTION: $ PUSH("GetExerciseInfo","FU_9"); $

STATE: FU_8
ACTION: $ if(!UNDEFINED(GET("COMMITMENT_EXERCISETYPE"))) {
    if (GET_EQ("COMMITMENT_EXERCISETYPE", "SPORT") &&
        !UNDEFINED(GET("COMMITMENT_SPORTTYPE"))) {
        GO("FU_12A");
    }
    else
        GO("FU_12");
}
else GO("FU_13"); $

STATE: FU_12A //known sport type
AGENT: $ <BEAT>Did you manage to play =|(GET("COMMITMENT_SPORTTYPE"))|= like you said you
would? </BEAT> $
REPEAT: $ <BEAT>Did you manage to play =|(GET("COMMITMENT_SPORTTYPE"))|= ? </BEAT> $

```

```

USERMENU:   Yes, according to plan. => $ GO("FU_9"); $
           No, I did something else. => $ GO("FU_7"); $

STATE: FU_12 //known exercise type
AGENT: $ <BEAT>Did you manage to =|EX_VP(GET("COMMITMENT_EXERCISETYPE"))|= like you said you
would? </BEAT> $
REPEAT: $ <BEAT>Did you manage to =|EX_VP(GET("COMMITMENT_EXERCISETYPE"))|=? </BEAT> $
USERMENU:   Yes, according to plan. => $ GO("FU_9"); $
           No, I did something else. => $ GO("FU_7"); $

STATE: FU_13 //unknown exercise type
AGENT: $ <BEAT>Did you manage to stick to your exercise plan? </BEAT> $
       $ <BEAT>Did you exercise as planned? </BEAT> $
       $ <BEAT>Did you follow your exercise plan? </BEAT> $
USERMENU:   Yes, according to plan. => $ GO("FU_9"); $
           No, I did something else. => $ GO("FU_7"); $

STATE: FU_9
ACTION: $ int actual=GETINT("EXERCISE_BEST_TIME");
         if(actual<GETINT("LAST_GOAL_TIME"))
           PUSH("RelapseTalk","FU_11");
         else PUSH("FollowUpSteps","FU_10"); $

STATE: FU_11
ACTION: $ PUSH("FollowUpSteps","FU_10"); $

STATE: FU_10
ACTION: $ int commitment=GETINT("COMMITMENT_TIME");
         int goal=GETINT("LAST_GOAL_TIME");
         int actual=GETINT("EXERCISE_BEST_TIME");
         boolean madeCommitment=!UNDEFINED(commitment);
         boolean didSomeExercise=(actual>=10);
         if(!didSomeExercise)
           PUSH("FollowUpObstacles","FU_END");
         else if(madeCommitment) {
           if(actual>=commitment)
             GO("FU_15");
           else if(actual<commitment && actual<goal)
             PUSH("FollowUpObstacles","FU_END");
           else
             GO("FU_16");
         } else { /* no commitment made */
           if(actual>=goal)
             GO("FU_17");
           else
             GO("FU_18");
         }
         };
$

STATE: FU_15
AGENT: $ <BEAT> You met your commitment, but. </BEAT> $
ACTION: $ PUSH("FollowUpObstacles","FU_END"); $

STATE: FU_16
AGENT: $ <BEAT> You met your goal, but didn't quite meet your commitment. </BEAT> $
ACTION: $ PUSH("FollowUpObstacles","FU_END"); $

STATE: FU_17
AGENT: $ <BEAT> You met your goal, but. </BEAT> $
ACTION: $ PUSH("FollowUpObstacles","FU_END"); $

//Would have already done RelapseTalk and said 'You didn't quite meet your exercise goals this
time.'
STATE: FU_18
ACTION: $ PUSH("FollowUpObstacles","FU_END"); $

STATE: FU_END
ACTION: $ POP(); $

```

"GetCommitment" script

```
STATE: GetCommitment
ACTION: $ if(!DEFINED(GET("TIME_TO_ASK_NEW_GOAL"))) PUSH("AskNewGoal","GC_1A");
        else GO("GC_1A"); $

STATE: GC_1A
AGENT: $ <BEAT><NEUTRAL/>So, let's talk about what you're going to do before we chat
again.</BEAT> $
        $ <BEAT><NEUTRAL/>So, let's chat about your near term exercise plans.</BEAT> $
REPEAT:$ <BEAT><NEUTRAL/>Let's chat about your near term exercise plans.</BEAT> $
ACTION: $ if(UNDEFINED(GET("GOAL_TIME"))) GO("GC_1X");
        else if(GETINT("GOAL_TIME")<GETINT("ULTIMATE_GOAL")) GO("GC_1B");
        else GO("GC_1C"); $

STATE: GC_1B
AGENT: $ <BEAT>Your long term goal is <GESTURE_RIGHT TYPE="CONTRAST_1" PRIORITY="10">
        =|GET("ULTIMATE_GOAL")|= minutes </GESTURE_RIGHT>
        of moderate or better activity a day.
        In order to work up to that you should currently be doing
        about <GESTURE_LEFT TYPE="CONTRAST_2" PRIORITY="10"> =|GET("GOAL_TIME")|= minutes
</GESTURE_LEFT>
        a day. </BEAT> $
ACTION: $ GO("GC_1X"); $

STATE: GC_1C
AGENT: $ <BEAT>You should be doing about =|GET("ULTIMATE_GOAL")|=
        minutes of moderate or better activity a day. </BEAT> $
ACTION: $ GO("GC_1X"); $

STATE: GC_1X
ACTION: $ SET("COMMITMENT_TIME","");
        SET("COMMITMENT_STEPS","");
        if(GET_EQ("NEXT_EXERCISE_DAY","TOMORROW"))
        GO("GC_16");
        else
        GO("GC_17");
        $

STATE: GC_16
AGENT: $ <BEAT>Are you going to work out tomorrow? </BEAT> $
USERMENU: I can't. => $ GO("GC_24"); $
        Yep. => $ GO("GC_18"); $

STATE: GC_24
ACTION: $ if(GET_SESSION_EQ("STATE","SICK") || GET_SESSION_EQ("STATE","INJURED"))
        GO("GC_20");
        else
        PUSH("MotivateToExercise","GC_33");
        $

STATE: GC_20
AGENT: $ <BEAT> Is it because of your
        =| ""+(GET_SESSION_EQ("STATE","SICK")? " illness? " : " injury?") |= </BEAT> $
USERMENU: Yes, it is. => $ GO("GC_23"); $
        No, not really. => $ PUSH("MotivateToExercise","GC_33"); $

STATE: GC_33 //check return value from MotivateToExercise
ACTION: $ if(GET_EQ("CURRENT_RETURN","TRUE")) GO("GC_18"); else GO("GC_34"); $

STATE: GC_34
AGENT_REL: $ <BEAT>OK. I trust you to make the right decision. </BEAT> $
AGENT_NREL: $ <BEAT>OK. I hope you're able to get back into the exercise hab it soon. </BEAT>
$
REPEAT: $ <BEAT>I hope you're able to get back into the exercise hab it soon. </BEAT> $
```

```

ACTION: $ GO("GC_END"); $

STATE: GC_23
AGENT: $ <BEAT>OK, I hope that you feel better soon, and can get back into the
exercise hab it. </BEAT> $ //tts
REPEAT: $ <BEAT>I hope that you feel better soon, and can get back into the
exercise hab it. </BEAT> $ //tts
ACTION: $ GO("GC_END"); $

STATE: GC_17
AGENT: $ <BEAT> Are you going to get some
=| (GETINT("EXERCISE_TIME") > 0? " more ": "") |= exercise today? </BEAT> $
USERMENU: Yes, I am. => $ GO("GC_18"); $
No, I'm not. => $ GO("GC_17A"); $

STATE: GC_17A
AGENT: $ <BEAT> Well, what about tomorrow? </BEAT> $
ACTION: $ GO("GC_16"); $

STATE: GC_18
AGENT: $ <BEAT> What kind of exercise are you going to do? </BEAT> $
USERMENU: I'm going to go for a walk. => $ SET("COMMITMENT_EXERCISETYPE", "WALK");
GO("GC_7"); $
I'm going to play a sport. => $ SET("COMMITMENT_EXERCISETYPE", "SPORT"); GO("GC_1");
$
I'm going to work out at the gym. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "GO"); GO("GC_3"); $
I'm going to run. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "RUN"); GO("GC_3"); $
I'm going to ride my bike. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "RIDE"); GO("GC_3"); $
I'm going swimming. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "SWIM"); GO("GC_3"); $
I'm going dancing. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "GO"); GO("GC_3"); $
Yoga. => $ SET("COMMITMENT_EXERCISETYPE", "GYM");
SET("CURRENT_V", "GO"); GO("GC_3"); $
I'm going to do something else. => $ SET("COMMITMENT_EXERCISETYPE", "OTHER");
GO("GC_19"); $

STATE: GC_7
ACTION: $ int goal=GETINT("GOAL_TIME");
if(goal<=30) GO("GC_7A");
else if(goal<=60) GO("GC_7B");
else GO("GC_7C"); $

STATE: GC_7A
AGENT_REL: $ <BEAT>Great. How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for?
</BEAT> $
AGENT_NREL: $ <BEAT>Great. Could you tell me how <INT_EMPHASIS/> long you plan to
<INT_EMPHASIS/> go for? </BEAT> $
REPEAT: $ <BEAT>How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for? </BEAT> $
USERMENU:
30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_32"); $
25 minutes => $ SET("COMMITMENT_TIME", "25"); GO("GC_32"); $
20 minutes => $ SET("COMMITMENT_TIME", "20"); GO("GC_32"); $
15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_32"); $
10 minutes => $ SET("COMMITMENT_TIME", "10"); GO("GC_32"); $
5 minutes => $ SET("COMMITMENT_TIME", "5"); GO("GC_32"); $

STATE: GC_7B
AGENT_REL: $ <BEAT>Great. How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for?
</BEAT> $
AGENT_NREL: $ <BEAT>Great. Could you tell me how <INT_EMPHASIS/> long you plan to
<INT_EMPHASIS/> go for? </BEAT> $
REPEAT: $ <BEAT>How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for? </BEAT> $
USERMENU:
75 minutes => $ SET("COMMITMENT_TIME", "75"); GO("GC_32"); $

```

```

        60 minutes => $ SET("COMMITMENT_TIME", "60"); GO("GC_32"); $
        55 minutes => $ SET("COMMITMENT_TIME", "55"); GO("GC_32"); $
        50 minutes => $ SET("COMMITMENT_TIME", "50"); GO("GC_32"); $
        45 minutes => $ SET("COMMITMENT_TIME", "45"); GO("GC_32"); $
        40 minutes => $ SET("COMMITMENT_TIME", "40"); GO("GC_32"); $
        35 minutes => $ SET("COMMITMENT_TIME", "35"); GO("GC_32"); $
        30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_32"); $
        15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_32"); $

STATE: GC_7C
AGENT_REL: $ <BEAT>Great. How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for?
</BEAT> $
AGENT_NREL: $ <BEAT>Great. Could you tell me how <INT_EMPHASIS/> long you plan to
<INT_EMPHASIS/> go for? </BEAT> $
REPEAT: $ <BEAT>How <INT_EMPHASIS/>long do you plan to <INT_EMPHASIS/> go for? </BEAT> $
USERMENU:
        90 minutes => $ SET("COMMITMENT_TIME", "90"); GO("GC_32"); $
        85 minutes => $ SET("COMMITMENT_TIME", "85"); GO("GC_32"); $
        80 minutes => $ SET("COMMITMENT_TIME", "80"); GO("GC_32"); $
        75 minutes => $ SET("COMMITMENT_TIME", "75"); GO("GC_32"); $
        70 minutes => $ SET("COMMITMENT_TIME", "70"); GO("GC_32"); $
        65 minutes => $ SET("COMMITMENT_TIME", "65"); GO("GC_32"); $
        60 minutes => $ SET("COMMITMENT_TIME", "60"); GO("GC_32"); $
        45 minutes => $ SET("COMMITMENT_TIME", "45"); GO("GC_32"); $
        30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_32"); $
        15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_32"); $

STATE: GC_32
ACTION: $
if(GETINT("COMMITMENT_TIME") < GOALTIME())
    GO("GC_9");
else if (GETINT("COMMITMENT_TIME")-GOALTIME())>=10) /* more than 10 mins over goal */
    GO("GC_21");
else
    GO("GC_8");
$

STATE: GC_9
AGENT: $ <BEAT> Do you think you can go for =| "+GOALTIME() |= minutes? </BEAT> $
        $ <BEAT> Could you try for =| "+GOALTIME() |= minutes instead? </BEAT> $
USERMENU: Sure.; I'll try. => $ SET("COMMITMENT_TIME",GET("GOAL_TIME")); GO("GC_8"); $
        No, I don't think so. => $ PUSH("MotivateDuration", "GC_31"); $

STATE: GC_21
AGENT: $ <BEAT>I think you shouldn't try to do so much so soon. How about just
        =| "+GOALTIME() |= minutes this time? </BEAT> $
REPEAT: $ <BEAT>How about just =| "+GOALTIME() |= minutes this time? </BEAT> $
USERMENU: Okay. => $ GO("GC_8"); $
        No, I really want to. => $GO("GC_22"); $

STATE: GC_22
AGENT: $ <BEAT>Okay, but remember it's better to increase your time gradually. </BEAT> $
ACTION: $GO("GC_8"); $

STATE: GC_8
ACTION: $
        if (ISREL() ) {
            if(UNDEFINED(GET("COMMITMENT_WALKLOCATION")))
                GO("GC_11");
            else
                GO("GC_10");
        }
        else
            GO("GC_31");
$

STATE: GC_10
AGENT: $ <BEAT>Are you going to walk

```

```

        =| "+"LOC_PP(GET("COMMITMENT_WALKLOCATION")) |=
        again? </BEAT> $
USERMENU: Yes, I am. => $ GO("GC_12"); $
        No, I'm going to a different place. => $ GO("GC_11");$

STATE: GC_11
AGENT: $ <BEAT><INT_EMPHASIS/>Where are you going to walk? </BEAT> $
        $ <BEAT><INT_EMPHASIS/>Where are you going? </BEAT> $
        $ <BEAT><INT_EMPHASIS/>Where will you go to walk? </BEAT> $
USERMENU:
In the park. => $ SET("COMMITMENT_WALKLOCATION", "PARK"); GO("GC_12"); $
By the water. => $ SET("COMMITMENT_WALKLOCATION", "WATER"); GO("GC_12"); $
In town. => $ SET("COMMITMENT_WALKLOCATION", "TOWN"); GO("GC_12"); $
On the treadmill. => $ SET("COMMITMENT_WALKLOCATION", "TREADMILL"); GO("GC_12"); $
At the running track. => $ SET("COMMITMENT_WALKLOCATION", "TRACK"); GO("GC_12"); $
Somewhere else. => $ SET("COMMITMENT_WALKLOCATION", "OTHER"); GO("GC_12");$

STATE: GC_12
AGENT: $ <BEAT> =| ("+(GET_EQ("COMMITMENT_WALKLOCATION",
"OTHER"))?":LOC_PP(GET("COMMITMENT_WALKLOCATION"))+"?") |=
        That sounds great. </BEAT> $
ACTION:
$
if (GET_EQ("ISLONER", "TRUE"))
    GO("GC_31");
else if (UNDEFINED(GET("BUDDY")))
    GO("GC_14");
else
    GO("GC_13");
$

STATE: GC_13
AGENT: $ <BEAT> Are <INT_EMPHASIS/>you going to go with <INT_EMPHASIS/>
        =| "+"((GET_EQ("BUDDY", "NOTCHECKED"))?" your buddy ":GET("BUDDY")) |=
        again? </BEAT> $
USERMENU: Yes, I am. => $ GO("GC_31"); $
        Not this time. => $ GO("GC_14"); $

STATE: GC_14
AGENT: $ <BEAT> Are you going to exercise <INT_EMPHASIS/>with anyone? </BEAT> $
USERMENU: Yes, I am. => $ GO("GC_15");$
        No, I'm going alone. => $ GO("GC_31");$

// note: if name is unknown to library, set name = "NOTCHECKED"
STATE: GC_15
AGENT: $ <BEAT> Terrific. <INT_EMPHASIS/>Who are you <INT_EMPHASIS/>going with? </BEAT> $
//tts
REPEAT: $ <BEAT> <INT_EMPHASIS/>Who are you <INT_EMPHASIS/>going with? </BEAT> $ //tts
USERTEXT: My walking buddy's name is: => $ SET("BUDDY",FILTERNAME(GETTEXT())); GO("GC_31"); $

STATE: GC_1
AGENT: $ <BEAT>Great. </BEAT> $
ACTION:
$
if (ISREL() && !UNDEFINED(GET("SPORTPLAYED")) && !GET_EQ("SPORTPLAYED", "OTHER"))
    GO("GC_2");
else
    GO("GC_4");
$

STATE: GC_2
AGENT: $ <BEAT> =| GET("SPORTPLAYED") |= again? </BEAT> $
USERMENU: Yep. => $ SET("COMMITMENT_SPORTTYPE", GET("SPORTPLAYED")); GO("GC_3"); $
        No, I'm doing another sport. => $ GO("GC_4");$

STATE: GC_4
AGENT: $ <BEAT><INT_SPEED WPM="140"/>Which <INT_EMPHASIS/> sport are you going to play?
</BEAT> $

```

```

USERMENU:
    I'm going to play squash => $ SET("COMMITMENT_SPORTTYPE", "SQUASH");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to play golf => $ SET("COMMITMENT_SPORTTYPE", "GOLF"); SET("CURRENT_V","PLAY");
GO("GC_3"); $
    I'm going to play tennis => $ SET("COMMITMENT_SPORTTYPE", "TENNIS");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to play hockey => $ SET("COMMITMENT_SPORTTYPE", "HOCKEY");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to play softball => $ SET("COMMITMENT_SPORTTYPE", "SOFTBALL");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to play basketball => $ SET("COMMITMENT_SPORTTYPE", "BASKETBALL");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to play soccer => $ SET("COMMITMENT_SPORTTYPE", "SOCCER");
SET("CURRENT_V","PLAY"); GO("GC_3"); $
    I'm going to do a different sport. => $ SET("CURRENT_V","PLAY"); GO("GC_3");$

```

```

STATE: GC_3
ACTION: $ int goal=GETINT("GOAL_TIME");
    if(goal<=30) GO("GC_3A");
    else if(goal<=60) GO("GC_3B");
    else GO("GC_3C"); $

```

```

STATE: GC_3A
AGENT: $ <BEAT>How long do you plan to <INT_EMPHASIS/> =|GET("CURRENT_V")|= for? </BEAT> $
USERMENU:
    30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_35"); $
    25 minutes => $ SET("COMMITMENT_TIME", "25"); GO("GC_35"); $
    20 minutes => $ SET("COMMITMENT_TIME", "20"); GO("GC_35"); $
    15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_35"); $
    10 minutes => $ SET("COMMITMENT_TIME", "10"); GO("GC_35"); $
    5 minutes => $ SET("COMMITMENT_TIME", "5"); GO("GC_35"); $

```

```

STATE: GC_3B
AGENT: $ <BEAT>How long do you plan to <INT_EMPHASIS/>play for? </BEAT> $
USERMENU:
    75 minutes => $ SET("COMMITMENT_TIME", "75"); GO("GC_35"); $
    60 minutes => $ SET("COMMITMENT_TIME", "60"); GO("GC_35"); $
    55 minutes => $ SET("COMMITMENT_TIME", "55"); GO("GC_35"); $
    50 minutes => $ SET("COMMITMENT_TIME", "50"); GO("GC_35"); $
    45 minutes => $ SET("COMMITMENT_TIME", "45"); GO("GC_35"); $
    40 minutes => $ SET("COMMITMENT_TIME", "40"); GO("GC_35"); $
    35 minutes => $ SET("COMMITMENT_TIME", "35"); GO("GC_35"); $
    30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_35"); $
    15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_35"); $

```

```

STATE: GC_3C
AGENT: $ <BEAT>How long do you plan to <INT_EMPHASIS/>play for? </BEAT> $
USERMENU:
    90 minutes => $ SET("COMMITMENT_TIME", "90"); GO("GC_35"); $
    85 minutes => $ SET("COMMITMENT_TIME", "85"); GO("GC_35"); $
    80 minutes => $ SET("COMMITMENT_TIME", "80"); GO("GC_35"); $
    75 minutes => $ SET("COMMITMENT_TIME", "75"); GO("GC_35"); $
    70 minutes => $ SET("COMMITMENT_TIME", "70"); GO("GC_35"); $
    65 minutes => $ SET("COMMITMENT_TIME", "65"); GO("GC_35"); $
    60 minutes => $ SET("COMMITMENT_TIME", "60"); GO("GC_35"); $
    45 minutes => $ SET("COMMITMENT_TIME", "45"); GO("GC_35"); $
    30 minutes => $ SET("COMMITMENT_TIME", "30"); GO("GC_35"); $
    15 minutes => $ SET("COMMITMENT_TIME", "15"); GO("GC_35"); $

```

```

STATE: GC_19
AGENT: $ <BEAT> What kind of exercise? </BEAT> $
USERTEXT: For exercise, I plan to: => $ SET("CURRENT_V","GO"); GO("GC_3"); $ //just ignore
the response...

```

```

STATE: GC_35
AGENT: $ <BEAT>So, you're going to exercise for =|GET("COMMITMENT_TIME")|= minutes. </BEAT> $

```

```

    $ <BEAT> =|GET("COMMITMENT_TIME")|= minutes? </BEAT> $
REPEAT:$ <BEAT> =|GET("COMMITMENT_TIME")|= minutes? </BEAT> $
ACTION: $ GO("GC_31"); $

STATE: GC_31
ACTION: $ if(UNDEFINED(GET("COMMITMENT_TIME"))) {
    if(GETINT("INTERACTION")>4) PUSH("GetStepsCommitment", "GC_30");
    else GO("GC_30");
} else if(GET_EQ("COMMITMENT_TIME","5")) GO("GC_31B");
    else if(GETINT("COMMITMENT_TIME")<GETINT("GOAL_TIME")) GO("GC_31D");
    else GO("GC_31C"); $

STATE: GC_31B
AGENT: $ <BEAT>Well, that's better than nothing. </BEAT> $
ACTION: $ if(GETINT("INTERACTION")>4) PUSH("GetStepsCommitment", "GC_30"); else GO("GC_30"); $

STATE: GC_31C
AGENT: $ <BEAT>That's great. </BEAT> $
ACTION: $ if(GETINT("INTERACTION")>4) PUSH("GetStepsCommitment", "GC_30"); else GO("GC_30"); $

STATE: GC_31D
AGENT: $ <BEAT>OK. </BEAT> $
ACTION: $ if(GETINT("INTERACTION")>4) PUSH("GetStepsCommitment", "GC_30"); else GO("GC_30"); $

STATE: GC_30
ACTION: $ if(!UNDEFINED(GET("COMMITMENT_TIME"))) PUSH("DoPerformative", "GC_END"); else
GO("GC_END"); $

STATE: GC_END
ACTION: $ POP();$

```

"Greeting" script

```

STATE: Greeting
ACTION: $ if(ISREL()) GO("GR_1"); else GO("GR_2"); $

STATE: GR_1
AGENT: $ <BEAT> <HAPPY/>Hi =|GET("NAME") |= </BEAT> $
    $ <BEAT> <HAPPY/>Hello, =|GET("NAME") |= </BEAT> $
    $ <BEAT> <HAPPY/> Hi =|GET("NAME") |=. Great to see you. </BEAT> $
    $ <BEAT> <HAPPY/>Hello, =|GET("NAME") |=. Great to see you. </BEAT> $
    $ <BEAT> <HAPPY/>Hey =|GET("NAME") |= </BEAT> $
    $ <BEAT> <HAPPY/>Good =|GET("PART_OF_DAY")|= , =|GET("NAME")|= ! </BEAT> $
    $ <BEAT> <HAPPY/>Hi, =|GET("NAME") |= I'm glad to see you. </BEAT> $
USERMENU: Hi Laura.; Hi Laura, great to see you. => $ GO("GR_END"); $
    Hi Laura, you don't look like things are great. => $ GO("GR_3"); $

STATE: GR_2
AGENT: $ <BEAT> Hi! </BEAT> $
    $ <BEAT> Hello. </BEAT> $
    $ <BEAT> Hello there. </BEAT> $
    $ <BEAT> Hey there. </BEAT> $
    $ <BEAT> Good <INT_EMPHASIS/> =|GET("PART_OF_DAY") |= </BEAT> $
USERMENU: Hello. => $ GO("GR_END"); $

STATE: GR_3
AGENT: $ <BEAT>Sorry, my animations are not rendering very smoothly today.</BEAT> $
    $ <BEAT>Sorry, I'm feeling a little rough around the edges.</BEAT> $
    $ <BEAT>Yeah, I'm having a pixelated hair day.</BEAT> $
ACTION: $ GO("GR_END"); $

STATE: GR_END
ACTION: $ POP(); $

```
