First Impressions in Human-Agent Virtual Encounters

Angelo Cafaro
Doctor of Philosophy
May 2014
School of Computer Science
Reykjavík University

Ph.D. DISSERTATION

ISSN 1670-8539

First Impressions in Human-Agent Virtual Encounters

by

Angelo Cafaro

Thesis submitted to the School of Computer Science
at Reykjavík University in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy

May 2014

Thesis Committee:

Hannes Högni Vilhjálmsson, Supervisor
Prof., Reykjavík University

Timothy W. Bickmore
Prof., Northeastern University

Dirk Heylen
Prof., University of Twente

Anna Esposito, Examiner
Prof., Second University of Naples

Copyright
Angelo Cafaro
May 2014

First Impressions in Human-Agent Virtual Encounters

Angelo Cafaro

May 2014

Abstract

Relational Agents (RAs) are virtual anthropomorphic characters able to engage in multimodal (i.e. using both verbal and nonverbal behavior) face-to-face interactions with users in real-time. RAs are also capable of establishing and maintaining a long-term relationship with the user, which has been shown to improve task outcomes in application domains such as education, coaching and entertainment. In all of these applications, it is crucial that users do not outright reject the agent after the initial moments of interaction, therefore first impressions become important. This thesis presents a theoretical framework for analyzing and modeling human nonverbal behavior for managing impressions, and how RAs can exploit this in their first encounters with human users. The focus is on nonverbal behavior (smile, gaze and proxemics) aimed at exhibiting personality and interpersonal attitudes. First the thesis describes the theoretical background of nonverbal communicative behavior in the context of first impressions among humans. It then presents a theoretical framework demonstrating that impressions of an agent's personality are quickly formed by users based on proxemics, whereas interpersonal attitude is conveyed through smiles and gaze behavior. The thesis furthermore demonstrates that interpersonal attitude has greater impact than personality on a user's decision to spend time with the agent. The design and implementation of a SAIBA compliant computational solution built on this framework is presented. This solution automates the real-time generation of nonverbal behavior for an RA during a greeting encounter. The multimodal behavior exhibited accomplishes both to serve the communicative functions associated with greetings as well as managing impressions of personality and interpersonal attitude towards the user. The agent's communicative functions are represented in Function Markup Language, an emerging SAIBA standard, which this thesis furthers through a detailed design specification and concrete language proposal. Finally, a practical application in the context of a 3D Virtual Learning Environment is demonstrated and the impact and further developments of this thesis discussed.

Áhrifarík fyrstu kynni af venslavitverum

Angelo Cafaro

Maí 2014

Útdráttur

Venslavitverur eru sýndarmanneskjur sem nýta tal og látbragð í samskiptum við fólk, augliti til auglitis, til að mynda varanleg tengsl. Sýnt hefur verið fram á að slík tengsl við notendur hugbúnaðar á sviði kennslu, þjálfunar eða afþreyingar, getur bætt áhrif hugbúnaðarins á notandann. Það er því mikilvægt að notendur hafni ekki vitverunni við fyrstu kynni og missi þar með af ávinningi tengslanna. Þess vegna er nauðsynlegt að veran komi vel fyrir strax í upphafi. Þessi ritgerð lýsir fræðilegum grunni sem nýtist til að greina og gera líkan af áhrifum látbragðs á þá mynd sem notendur gera sér af vitverunni, og þar með hvernig vitveran getur best beitt þeirri tjáningu til að tryggja að á komist gott samband. Sérstök áhersla er lögð á látbragð sem gefur í skyn ákveðinn persónuleika og viðhorf gagnvart viðmælanda. Fyrst lýsir ritgerðin því hvernig fræðin taka á þessu þegar aðeins manneskjur eiga í hlut. Síðan greinir ritgerðin frá rannsókn sem sýnir hvernig líkamleg nálegð vitveru hefur áhrif á mat notanda á persónleika hennar og hvernig bros og augnaráð verunnar hefur áhrif á mat notanda á viðhorfi hennar til hans. Jafnframt sýnir ritgerðin að metið viðhorf hefur meiri áhrif en metinn persónuleiki á það hversu miklum tíma notandi er tilbúinn að eyða með vitverunni. Þessar niðurstöður eru útfærðar í reiknilíkani sem fellur að hinni alþjóðlegu SAIBA aðferð við myndun marghátta samksiptahegðunar hjá vitverum. Líkanið framkallar sjálfvirk alla líkamstjáningu vitverunar þegar hún og notandinn heilsast. Líkamstjáningin þjónar bæði athöfninni að heilsast og því að koma ákveðnum persónuleika og viðhorfi til skila. Samskiptamarkmið vitverunnar eru sett fram á "Functional Markup Language" málinu, sem er SAIBA staðall í þróun. Með þessu leggur ritgerðin til nákvæma greiningu og tillögu að fyrstu heildarútgáfu staðalsins. Að lokum lýsir ritgerðin notkun líkansins í gagnvirku þrívíddar kennsluumhverfi ásamt því að fjalla um næstu skref verkefnisins og áhrif þess á frekari rannsóknir á sviðinu.

Dedicated to my family,
who provided support along every step of the way.

Acknowledgements

It is hard to describe with words my gratitude for Hannes Högni Vilhjálmsson, the supervisor of this work. He thought me the art of science by the means of a genuine passion for his work and a truly unique good temper. He always had broad visions, but at the same time he was capable of narrowing his supervision down to the smallest details. He was constantly present but never too invasive or authoritative, he made sure there was always room for my imagination and creativity. Thanks to him I appreciated the benefits of being open minded and respectful towards others' work. I learned the importance of assuming a positive attitude also in the dark moments. I left behind my fears and concerns, and I feel ready to afford all the challenges that lie ahead in both my professional and personal life.

Profound gratitude goes to his wife Deepa Iyengar. She was supportive and showed interest for our work in every occasion. I admire her personality and sense of humor.

I would like to thank my committee members Timothy Bickmore and Dirk Heylen for believing in this work. They always provided me with thought provoking comments and showed the importance of considering my own work from different perspectives. Special thanks to my thesis readers and the external examiner, Anna Esposito, for their great feedback.

I would like to express my gratitude to our three main collaborators: Gunnar Steinn Valgarðsson, Daniel Schulman and Kamilla Rún Jóhannsdóttir. I am grateful to Gunnar for lending me the virtual character that he created allowing me to quickly set up and run our research studies. He was my trusted technical consultant with practical suggestions for every issue I presented to him. I am indebted to Daniel for sharing his knowledge in statistics and for providing great feedback on my work. Following his suggestions I was never wrong at applying the cor-

rect statistical methods and he sparked in me the interest for such a complex but beautiful field. Thanks to Camilla for inspiring discussions and for sharing her expertise in cognitive psychology to design our empirical experiments.

Big thanks to the Relational Agents Group at Northeastern University and the staff members at the Museum of Science in Boston. My visit for a semester there under the guidance of Timothy Bickmore deeply contributed to my personal and academic growth. I was lucky to work with an amazing group. I got endless inspiration from the enthusiasm and intellectual curiosity of Barbara Berry. I got incredible support from Juan Fernandez on all the trials and tribulations with Tinker. I was impressed by the patience of Lazlo Ring in answering every question I ever had while working together. Thanks to Andrea Tommaso Bonanno and Langxuan (James) Yin for cheering me up all the times and showing me around all the secrets of beautiful Boston.

Thanks to the Center for Analysis and Design of Intelligent Agents (CADIA), the School of Computer Science at Reykjavik University, the Icelandic Research Found (RANNIS) and the National Science Foundation (NSF) in USA for their financial support.

Special thanks to all past and present CADIA members and visiting students. In particular, I would like to thank Alfredo Vincenzo Conte, Andrea Monacchi, Carmine Oliva, Claudio Pedica, Elín Carstensdóttir, Fabrizio Fornari, Francesco Paolo Pandolfo and Lorenzo Scagnetti for the fruitful exchange of feedbacks we had from time to time and for sharing with me the lab life.

I would like to thank Reykjavik University for offering access to great facilities and a formidable set of staff members and professors. I am indebted to the IT services for their technical support while running our research studies. I appreciated a lot the efficiency of the library, the travel services and the student office. I had the opportunity to interact with experienced professors, in particular with Luca Aceto and Magnús Már Halldórsson, and receive from them decisive intellectual aid on some hard theoretical problems that I had to face with.

There is a “gang of dudes” that made my experience in Iceland unforgettable. The legendary members of this gang are: Claudio Pedica, Eugen Ioan Goriac, Georgiana Caltais, Matteo Cimini, Niccolò Giulio Rossetti, Ólafur Baldvin Jónsson, Pradipta Prometheus Mitra and Victor Carazo Robles. Their friendship is invaluable and so worth it that I would start over a thousand times my PhD just to experience again with them all the beautiful moments that we had.

“The Gang”. In the photo on the left from left to right: Niccolò Giulio Rossetti, Georgiana Caltais, Angelo Cafaro, Eugen Ioan Goriac, Matteo Cimini, Claudio Pedica and Pradipta Prometheus Mitra. In the photo on the right from left to right: Ólafur Baldvin Jónsson, Claudio Pedica, Angelo Cafaro and Victor Carazo Robles.

Thanks to these wonderful persons I’ve lived plenty of good times and never ending stories to tell.

Deepest gratitude and love goes in particular to Victor. He is the best person I have ever met so far at donating his time to friends. It is impractical to describe here what this guy has done for me. I consider him a brother, and I really mean it.

Sincere esteem for Guðmundur Albert Birgisson and Unnur Jóhannsdóttir for renting me out their apartment in the heart of Reykjavik, the legendary place at Hofsvallagata, 1. It was my safe harbor in Iceland and their daughter Edda Sonja Guðmundsdóttir and them were my little family there.

I am grateful to Alessandra De Martino, Davide Somma, Enrico Rossomando, Enzo Lo Sasso, Marco Alfano, Pasquale Picarella, Pietro Rossini and Raffaele Gaito for taking their time and coming from Italy to visit me. Thanks to Ásta Kristín Benediktsdóttir for recommending me beautiful trips around Iceland and for the interesting chats that we had from time to time in the cozy coffee shops downtown Reykjavik. I am indebted to Kristinn Már Jóhannesson for his support in the very first months of my stay and to Alessandra Procopio, Grazia Russo, Guðrún Kristín Einarsdóttir and Joseph Lovecchio for cheering me up in the very last months, when this final dissertation was only a bubble in my mind.

Warm thanks to my beloved family. I wish I had visited them more often, they have been present and supportive along every step of the way.

My warmest thanks to Vincenza Rossini. Words have come to an end here, space for feelings is needed.

Contents

List of Figures

List of Tables

"Nothing is more scientific than imagination."

Roberto Benigni (1952 – present)

1

Introduction

1.1 The Focus of This Thesis

Embodied Conversational Agents (ECAs) are anthropomorphic interface agents able to engage a user in real-time, multimodal dialogue, using speech, gesture, gaze, posture, facial expressions, intonation and other verbal and nonverbal channels to emulate the experience of human face-to-face interaction [Cassell et al., 2000b].

In educational [Vala et al., 2007], entertainment [Hartholt et al., 2009], training [Johnson et al., 2004] applications or as natural human-computer interfaces [Hoque et al., 2013, Cassell et al., 1999], ECAs are an important mechanism for delivering a variety of content and scenarios (see examples in Fig. 1.1). Already from the very first moment of interaction with the user, it is fundamental that the behavior exhibited by an agent is believable for the given context and content of the scene including other characters and the user. This is challenging because people generally expect a virtual agent to behave in a manner befitting its appearance and will often be disturbed by discrepancies in its behavior [Vinayagamoorthy et al., 2006].

(a) MACH (b) Icelandic Training

Figure 1.1: Two examples of ECA applications: (a) MACH, My Automated Conversation coacH, is a system for people to practice social interactions in face-to-face scenarios. As a paradigm of natural human-computer interaction, the system provides a virtual agent that reads users' facial expressions, speech, and prosody and responds with verbal and nonverbal behaviors in real time in the context of training for job interviews [Hoque et al., 2013]. (b) In the Icelandic Language and Culture Training system learners acquire basic communicative skills in Icelandic engaging in simulated social interactions with ECAs in a 3D virtual environment featuring downtown Reykjavik, in Iceland.

The challenges increase when those interface agents are built to carry long-term interactions with their users, share trust, feel empathetic, and effectively establish a connection, a bond. Those are known as Relational Agents (RAs) [Bickmore et al., 2005] and being able to build a socio-emotional relationship with their users has been shown to improve task outcomes in application domains such as coaching [Watson et al., 2012], counseling [Kang et al., 2012], psychotherapy [Rizzo et al., 2013] and health care [Edwards et al., ress].

Every interaction we have with a new person entails a first impression [Nalini and Skowronski, 2008], likewise for a user it might be desirable to encounter agents that are truly alive and approachable, in order to be encouraged to engage in an interaction and have the opportunity to befriend the agents. Obviously, if a user avoids future contact with an agent, his or her impression of this agent is unlikely to change, therefore the challenge is to ensure that subsequent interaction will even occur by making sure that the users do not outright reject those agents at the first encounter.

Under the fundamental assumption that socio-psychological principles of human-human interaction can be applied to human-agent interactions, we propose a theoretical framework to analyze, understand and model human nonverbal communicative behavior exhibited by a relational agent during the greeting phase of a first virtual encounter with the user. The goal is to develop a new computational solution based on the proposed framework for building more believable and effective RAs capable of managing the first impressions users have of them.

We demonstrate an application of our component in the context of Virtual Learning Environments (VLEs). In particular, we integrate our solution with the "Icelandic Language and Culture Training in Virtual Reykjavik" project[1], where learners interact with various autonomous inhabitants of Reykjavik to learn basic communicative skills in Icelandic, while under the guidance of a relational tutoring agent.

1.2 A First Virtual Encounter with a Relational Agent

Consider the language training environment mentioned above. In this environment you can navigate through a virtual reconstruction of the city populated by autonomous human-like virtual agents with whom you can interact with and train your skills. Some of these agents in the simulated environment are just passers-by, but others are special. They are meant to be your fellow adventurers, your tutors.

You've just got familiar with the dynamics of the system and you're going to start you first task, that consists of meeting your tutor and start your first training session. With your avatar (user controller character) you might try to get closer to be able to interact, but you realize that the agent does not show any reaction to your presence. You gradually approach him hoping to be seen, but it seems to be disinterested in you, almost distracted or maybe angry, a bit sad, or just simply not a sociable "person". Then all of a sudden, while you're less than a meter away, it suddenly smiles and looks at you.

What do you think of this agent? Does it really matter for you how this agent could be? Would you go ahead and start a training session with it? Assuming that you disliked it, maybe the agent was just programmed that way to react to your presence, or maybe it was a design mistake. Would you give it a second chance the next time you train if it is still there?

The tutors in this example are relational agents. Though they are designed to support long-term interactions with their users, in those very first steps of human-computer interaction, given their embodiment capabilities, they face an important choice:

---

[1] More details to follow at the web page: `http://secom.ru.is/`

How to select the right nonverbal behavior cues when approached by a user to avoid making an unwanted impression?

The scenario we are focusing on is a stranger-to-stranger situation, also known as a *zero-acquaintance* situation in social psychology literature [Nalini and Skowronski, 2008, p. 129]. While in this context most of the users' judgments can be made during the face-to-face conversation with an agent, during the very first moments of a virtual greeting encounter, even before the conversation takes place, there might be some nonverbal cues exhibited by the agent (e.g. a smile or some eye contact with the user) that assume a fundamental role for users' first impressions in such context.

Greetings have an important function in the management of relations, for example to begin, confirm and continue a friendship [Kendon, 1990, p. 154]. Nonverbal reactions, such as smile, gaze and proxemics behavior become powerful signals, indicating the desire to interact[Argyle, 1988]. Smiling is a common feature of greeting rituals [Kendon, 1990, p. 189]. Spatial moves (proxemics) are also very powerful means for starting and ending an encounter [Argyle, 1988], though they can be accompanied by other nonverbal signals, such as looking [Kendon, 1990, p. 189] and smiling at the person [Argyle, 1988, p. 178]. Spatial behavior may also signal a particular definition of the situation. Just as orientation communicates co-operative and competitive relationships, so a greater distance indicates the desire for greater formality [Argyle, 1988].

Because of their presumed centrality to relational communication, the specific subset of cues identified by smiling, gazing and proxemics is also categorized as "nonverbal immediacy behavior". They function to increase or decrease the physical and psychological distance between people, but they can also be a valuable means for determining the personality or attitude of a communicator [Richmond et al., 2008].

During even the most fleeting interactions, perceivers rapidly form impressions of one another's personality traits [Kammrath et al., 2007]. Those impressions can be surprisingly accurate - expressive behavior reveals a great deal even in small doses [Nalini and Skowronski, 2008, p. 111]. This is fundamental for the consistency of our relational agents. Guidelines for creating characters often include a caveat that everything a character does should convey the same general impression about the character to the viewer [Thomas and Johnston, 1981]. In particular, developing a clear, consistent, and appealing "personality" is an important part of creating successful characters [Isbister and Nass, 2000].

Interpersonal attitudes of a person toward another can vary in different contexts and have shorter duration. Sometimes people are trying to establish a relationship with another person, for example, being friendly or showing politeness during first encounters [Richmond et al., 2008]. As with expression of personality, there are spontaneous and managed aspects of interpersonal attitudes that can be expressed by nonverbal behavior [Argyle, 1988] and need to be addressed when approaching a relational agent. For example are you more likely to approach someone who is friendly towards you?

In addition to helping us managing social relationships while interacting with others, and "giving off" impressions of personality and attitude during greeting encounters. Nonverbal behavior plays a vital part in regulating the interaction, for example allowing people to take turns in a conversation seamlessly using gaze [Argyle and Cook, 1976, Argyle et al., 1973], and it provides information to a conversation (e.g. illustrative gestures [McNeill, 1992]).

Thus, different nonverbal behavior choices contribute to different and fundamental communicative functions that are extremely important to understand [Vilhjálmsson, 2009]. In the context of a first encounter, we aim to understand how the greeting communicative functions play with each other. In particular, the combination of immediacy cues described, first, it might be misinterpreted when blending those cues together, considering that they account for impressions of both personality and attitude at the same time, but they also allow someone to accomplish basic communicative functions such as greeting or inviting to interact. Secondly, other kind of unwanted judgments might arise when users will interact with our agents.

1.3 Approach

This thesis presents a new way to enhance relational agents' design by making them more effective from the very first moment of interaction with the user. We aim for the "right choice" of nonverbal immediacy behavior at the "right moment" of interaction, that is during the approach of the user towards the agent in a first greeting encounter.

When it comes to interactions with RAs, it is not so obvious whether the theories briefly illustrated above (and detailed in chapter 2) are still valid and can be trusted blindly. Our approach combines models drawn from the body of existing literature in sociology and psychology with new gathered data from user controlled experiments to attain a theoretical framework aimed at understanding:

1. How users interpret the nonverbal behavior of agents in terms of impressions of personality and attitude;
2. The impact of these impressions on users' relational decisions in terms of how likely they would be to interact with the agent again and how often they would interact with it;
3. The effects of managing first impressions of a relational agent in a real application setting.

Addressing this last point allows us to test the effectiveness of our approach in a different application domain. We believe that engaging users and making favorable impressions is a primary requisite for real life deployments of relational agents too. For example in applications, such as the museum installations described in [Bickmore et al., 2013, Kopp et al., 2005], that might have thousands of people passing near the agent and quickly deciding whether or not to approach it.

One complication in addressing these challenges is that studies in interpersonal psychology have shown that people tend to prefer others based on the match or mismatch to their own personality [Nass et al., 2000]. Sociable people are more highly motivated to detect and utilize behavioral cues to social relations, and are able to form impressions of others' personality traits with greater accuracy when interpreting observed nonverbal cues in first encounters [Nalini and Skowronski, 2008, p. 94,108,121]. A person's disposition can influence his or her evaluation and actual behavior while interacting with ECAs as well. An experiment with a real estate relational agent (REA) [Bickmore and Cassell, 2001], for example, showed that the agent's use of small talk increased trust in it for extroverts, but for introverts it had no effect. So the effectiveness of an agent depends on user's personality too [Pütten et al., 2010, Kang et al., 2008a] and we take this aspect into account in our framework.

The theoretical framework will be the foundation for a software component aimed at enhancing RAs greeting capabilities by automating the nonverbal behavior selection, including their positional and orientational parameters, to convey specific impressions of personality and attitudes toward the user. In order to untie this component from the specific application domain of VLEs and become a general-purpose module usable in different application domains and RAs implementations, we rely on the SAIBA framework (described in more detail in chapter 5). SAIBA is a standard behavior generation framework for ECAs that allows researchers to pool their efforts and speed up construction of multimodal interaction systems. For our work we propose a new standard specification for communicative functions (FML: Function Markup Language) to implement an easy to plug-in solution independent of the way the agent is rendered (2D or 3D) and the mode the user interacts with it (physical approach or avatar based interaction).

1.4 Contributions and Organization of Thesis

First the **theoretical background** about nonverbal communicative behavior in the context of first impressions in human social interactions is reviewed in chapter 2 along with details about the theoretical models that we adopted as basis of our work. Then we provide a review of **related works** in the field of embodied conversational agents and relational agents in chapter 3. In chapter 4 we present a **theoretical framework** for analyzing users' first impressions of an agent based on a series of empirical studies and a model of human nonverbal communicative behavior in greeting encounters. In chapters 5 and 6 we move from theory to practice. First, we discuss the design principles and propose the specification for a new standard to represent communicative functions in multimodal communication, the **Function Markup Language** (**FML**), in chapter 5. Then, in chapter 6, we apply our theoretical framework and FML to the design and implementation of a **computational solution** that automates the generation of an agent's nonverbal communicative behavior in virtual greeting encounters. A **working application** of this component in the context of Virtual Learning Environments will be also demonstrated. In chapter 7 we discuss the impact of our work, interesting issues that have been addressed as well as limitations of our approach. Finally, in chapter 8, we show where this research can lead us, including new interesting research questions, important considerations to take into account for the next steps and possible follow-up studies.

"There can be as much value in the blink of an eye as in months of rational analysis."

Malcolm Gladwell - Blink: The Power of Thinking without Thinking

2

Theoretical Background

2.1 First Impressions

Erving Goffman in his book "The Presentation of Self in Everyday Life", uses the imagery of theater in order to portray the importance of human and social action and interaction. He refers to this as the dramaturgical model of social life. According to Goffman, people in everyday life represents actors on a stage, each playing a variety of roles. The audience consists of other individuals who observe the role-playing and react to the performance [Goffman, 1959]. Central to Goffman's concern is his notion of *impression management* and *formation*:

Impression management Also called self-presentation, refers to the process by which individuals attempt to control the impressions that others form of them. Because the impressions people make on others have implications for how others perceive, evaluate, and treat them, as well as for their own views of themselves, people whether or not thinking about it, are often engaging in impression management, trying to control the information that others receive about them [Miller et al., 2007, Leary and Kowalski, 1990] to convey particular impressions [Goffman, 1959].

Impression formation It is the process by which individuals perceive, organize, and ultimately integrate information to form unified and coherent situated impressions of others. Internalized expectations for situated events condition what information individuals deem is important and worthy of their attention. Further, these expectations condition how individuals interpret this information and serve as the basis for subsequent attributions [Moore, 2006].

The "setting" for the performance includes the scenery, props, and location in which the interaction takes place. Different settings will have different audiences and will thus require the actor to alter his performances for each setting. The concepts of impression management and formation are, however, equally applicable to the different contexts where first impressions arise, some examples are: first encounters [Ambady and Rosenthal, 1993], job interviews [Stevens and Kristof, 1995], sporting encounters [Greenlees et al., 2005] and even on web pages [Weisbuch et al., 2009]. For this last example, it is notable to see that researchers found correlations between impressions formed from face-to-face interaction and personal web pages. People liked by interaction partners were also liked on the basis of their Facebook pages!

Goffman uses the term "performance" to refer to all the activity of an individual in front of a particular set of observers, or audience. Through this performance, the individual, or actor, gives meaning to themselves, to others, and to their situation. These performances deliver impressions to others and information is exchanged to confirm identity. The actor may or may not be aware of their performance or have an objective of their performance, however the audience is always attributing meaning it and to the actor. Therefore, from a communicative point of view, it should also be clear that the processes of impressions management and formation always co-exists.

Following this theatrical analogy, our relational agents perform as actors when interacting with their audience, that consists of the users. Users might not merely form impressions of those agents based on their "performance", but they might also undertake important decisions about a potential future relationship with them, according to the *Predicted Outcome Value* theory.

2.2 Predicted Outcome Value Theory

Consider a brief first encounter such as a job interview, a blind date or the first meeting with a potential roommate, the judgments we form in all of these situations often have enormous staying power, with our initial perceptions influencing us for months [Miller et al., 2007, Nalini and Skowronski, 2008] and determining important relational decisions such as the likelihood and frequency of subsequent encounters [Riggio and Friedman, 1986]. Even when our impressions of others turn out to be inaccurate, they can be powerful determinants of how we behave towards them [Rule and Ambady, 2010a] and, importantly, they shape expectations that we bring to future encounters [Goffman, 1959]. A work of [Sunnafrank and Ramirez, 2004] showed that relational partners make decisions within the first few weeks of a relationship that determine the long-term nature of the relationship. The theoretical background under this assumption is the *Predicted Outcome Theory* (POV). POV proposes that in social environments where potential relational partners will be proximate to one another in the future, individuals assess the likely outcomes of a future relationship to determine whether to develop the relationship and, if so, what type of relationship to attempt and how to proceed [Sunnafrank, 1986]. We gather information about others and our potential relationship through a combination of observation and interaction strategies [Berger and Kellerman, 1994]. When successful, the acquired information results in a more certain impression of others and the relationship potential [Sunnafrank and Ramirez, 2004].

The predicted outcome assessed by users in a first virtual encounter with an agent might be tremendously impactful on their relational decisions determining whether or not they will engage in subsequent interactions.

Furthermore, evidence of a connection between first impressions and nonverbal behaviors has been widely demonstrated [Argyle, 1988, Nalini and Skowronski, 2008, Richmond et al., 2008]. Before we speak, our gestures, postures and facial expressions are already broadcasting messages to those around us [Marsh and Gilmour, 1988] and we use all available information, particularly in the form of facial displays, gestures and other nonverbal cues, to infer states and traits that are intentionally or not directly sent to us [Nalini and Skowronski, 2008, p. 120].

2.3 The Role of Nonverbal Behavior

First impressions in greeting encounters are shaped both by individual characteristics or stereotypes that stay *invariant* over the interaction time, such as skin color, height, clothing or, generally, visual appearance [Bar et al., 2006, Naumann et al., 2009, Miller et al., 2007, Argyle, 1988] and by *dynamic* characteristics, such as verbal [Leary and Kowalski, 1990, Stevens and Kristof, 1995] and nonverbal behavior [Mehrabian, 1969, Riggio and Friedman, 1986, Burgoon et al., 1984, Argyle, 1988], that have higher degree of variability. In general, we have a different level of control among those characteristics. Individuals, for example, can carefully plan how to visually present themselves in a first encounter and what to say, but then it may be difficult to have full control over all nonverbal cues during the interaction [DePaulo, 1992]. In fact, one of the most interesting properties of nonverbal cues in social interaction is that they are irrepressibly impactful. Try as they might, people cannot refrain from behaving nonverbally. If, for example, they try to be as passive as possible, they are likely to be perceived as unexpressive, inhibited, withdrawn, and uptight [DePaulo, 1992]. Therefore, nonverbal behavior plays a fundamental but, at the same time, subtle role in the dynamics of impression management.

Goffman, in his analogy, refers to verbal and nonverbal behavior as "sign vehicles" that might arise from two different kinds of expressions:

1. The expressions "*we give*": primarily the things we say, intentional poses, facial expressions (smiles, surprise, etc.) and other controlled body language we emit;
2. The expressions "*we give off*": the elements of our expressiveness over which we have less control. The inconsistencies between what we say and what we actually do, the body language which "gives us away".

Either intended or unintentional, our first impressions of others can be truthful [Bar et al., 2006]. Researchers have demonstrated that people accurately infer a remarkable amount of information including, for instance, a person's skill level [Ambady and Rosenthal, 1993, Rule and Ambady, 2010b], personality [Campbell and Rushton, 1978, Levesque and Kenny, 1993], sexual orientation [Rule and Ambady, 2008, Rule et al., 2009], political view [Rule and Ambady, 2010c], religious group membership [Rule et al., 2010], later career success [Rule and Ambady, 2011], race or age [MacLin and Malpass, 2001], after observing someone's nonverbal behavior for a relatively short period of time (for example 2s, 5s, 10s video clips or photos). Interestingly, in the case of video stimuli, longer slices are not associated with significant increases in predictive accuracy. [Ambady and Rosenthal, 1993] found that students, for example, were surprisingly good at predicting a teacher's effectiveness based on first impressions. They videotaped 13 graduate teaching fellows as they taught their classes and, then, took three random 10-second clips (termed "thin slices") from each tape, combined them into one 30-second clip for each teacher and showed the silent clips to students who did not know the teachers. The student judges rated the teachers on 13 variables, such as "accepting", "active", "competent" and "confident". Finally, they combined these individual scores into one global rating for each teacher and then correlated that rating with the teachers' end-of-semester evaluations from actual students, resulting in a very high correlation.

Relational agents compared to humans can benefit from not having to deal with this intended or unintentional distinction. When it comes to human-like appearance, the artificial nature of this embodiment allows them to have full control on their realization capabilities. This solves the problem of unwanted nonverbal displays, however choosing the right nonverbal behavior still remains a challenge, since it often carries a communicative meaning that cannot be neglected if we want to avoid miscommunication with the user.

2.4 Nonverbal Communicative Functions and Immediacy

Communicative Functions in Multimodal Communication

Nonverbal behaviors, contribute to fundamental communicative functions that are extremely important to understand when modeling a virtual anthropomorphic agent. These functions can be broadly characterized in three main categories [Vilhjálmsson, 2009]:

- **Interaction Functions.** This category has to do with regulation and effective management of an interaction. For example nonverbal displays of interest often determine whether or not interaction ever begins, and, thereafter, subtle nonverbal cues allow people to take turns in a conversation seamlessly and gracefully;
- **Content Functions.** This one covers the actual content that gets exchanged across live communication channels. The various functions belonging to this category can be divided across different organizational levels, from the largest organizational structure of discourse (e.g. topics and segments) and information structure (e.g. new or given content) down to the specification of each proposition;
- **Mental States and Attitude Functions.** This last category deals with functions describing cognitive processes (i.e. planning, thinking and remembering), emotional states and socio-relational activities (i.e. expressing intimacy or affiliation, carrying out signals of power and status).

In our work the link between nonverbal behaviors and the last category of functions, accomplishing socio-relational activity goals, is particularly interesting. Nonverbal cues play a central role in personal and social perception [Burgoon and Le Poire, 1999] and are especially crucial for establishing and building a relationship, since they can be used to provide such social cues as attentiveness, positive affect, liking and attraction, and to mark shifts into and out of relational activities [Bickmore and Cassell, 2001]. [Burgoon et al., 1984] showed the association between patterns of nonverbal behaviors and relational messages such as detachment, emotional arousal, dominance and control. Judgment of trustworthiness, reliability and higher likelihood of forming working alliances are related to smiling [Mehu et al., 2008]. Gaze also helps define the relationship two people share once the interaction begins [Miller et al., 2007].

Nonverbal Immediacy

According to [Mehrabian, 1969], body posture, distance and body orientation, eye contact and smiling, also categorized as “immediacy behavior” [Richmond et al., 2008], have been found most consistently to be indicators of communicator attitude toward an addressee, where attitude is broadly defined as the degree of liking, positive evaluation, and/or preference of one person to another.

Immediacy can be defined as the degree of perceived physical or psychological closeness between people [Richmond et al., 2008]. Nonverbal immediacy refers to the use of nonverbal behavior that increases the immediacy between interactants. The more forms of approaching behavior we use, the more we are perceived as non-verbally immediate, and this in turn might have positive outcomes for the purpose of building relationships: standing close to and leaning forward another person, smiling, making eye contact, facing the other person appropriately allow one to make a favorable impression on that person [Miller et al., 2007, Richmond et al., 2008].

The Importance of Smile, Gaze and Proxemics cues

Smiling behavior, for example, could considerably help the process of building social relationships [Mehu et al., 2008] and has been related to having positive impact on a measure of interpersonal attraction represented by willingness to work together, friendliness and approachability [Mehu et al., 2008]. **Posture** creates an impression of attitude, some body shapes create a stronger impression of interest, agreement [Marsh and Gilmour, 1988], and interpersonal attitudes (friendliness, hostility, defensiveness, etc. . .) [Mehrabian, 1969] than other shapes for example. Patterns of **eye contact** can be influenced by some personality factors [Marsh and Gilmour, 1988], for examples those who have an outgoing and sociable personality, look at other people more than introverts do [Argyle, 1988].

According to Argyle, **distance** (or **proxemics** behavior) and **body orientation** belong to a broader category of behaviors named **spatial behaviors** [Argyle, 1988]. They are one of the main ways to express friendly-hostile attitudes to other people and, in some way, they are the most straightforward nonverbal signals, since they can be easily measured in terms of distance or orientation. Interpersonal distance expresses certain personality qualities as well. Patterson and Schrest [Argyle, 1988, p. 178] asked subjects to interview a number of people who had been instructed to approach to different distances. They sat on one of a series of chairs placed side by side at different distances. The impression formed of these people varied with the distance they chose. Those who came closest were rated as friendly, extroverted, aggressive and dominant.

In conclusion, by properly selecting nonverbal immediacy cues consisting of **smiling**, **gazing** or **proxemics** behavior, it is possible to increase or decrease the physical and psychological distance between people [Burgoon et al., 1984], and they can be a valuable means for determining the personality [Campbell and Rushton, 1978] or attitude of a communicator [Argyle, 1988, Mehrabian, 1969]. The ability to gauge an individual's personality and attitude quickly and on the basis of limited information is of critical importance when considering starting a relationship and forming those impressions occurs spontaneously, without intention or even awareness [Nalini and Skowronski, 2008, p. 108].

2.5 Personality Traits and Interpersonal Attitudes

Big 5 Personality Traits Model

Personality researchers have identified a major approach to the study of human personality by measuring traits, which can be defined as broad themes in behavior, thoughts, and emotions that distinguish one person from another [Miller et al., 2007]. The traits used to describe personality ranges from a few to, potentially, an unlimited number. Hans Eysenck, for example, has suggested that personality is reducible to three major traits [Eysenck, 1991], whereas other researchers argue that more factors are needed to adequately describe human personality [Goldberg and Saucier, 1998]. Concerted efforts at factors analysis, however, have demonstrated that five factors are sufficient for providing the best compromise between explanatory power and parsimony. These are the "Big Five" personality traits [McCrae and Costa, 1997] and can be summarized as in [Miller et al., 2007]:

- **Extraversion**: The extent to which people are outgoing, gregarious, talkative, and sociable versus cautious, reclusive, and shy;
- **Agreeableness**: The degree to which people are good-natured, cooperative, and trusting versus irritable, cranky, and hostile;
- **Conscientiousness**: The extent to which people are responsible, dutiful, and dependable versus unreliable and careless;
- **Neuroticism**: The degree to which people are impulsive and prone to worry, anxiety, and anger;
- **Openness to experience**: The degree to which people are imaginative, unconventional, and artistic versus conforming, uncreative, and stodgy.

The table below shows a few traits that represent high and low values along each dimension:

Some of those traits are more visible than others [Miller et al., 2007, p. 140-141]. Extraversion, in particular, seems to be one of the easiest trait to pick [11,18] and could represent crucial information to answer adaptive questions such as whom to mate with or rely on in social alliances [Mehu et al., 2008]. With only a brief glimpse at expressive behavior, including interpersonal distance, smile, gaze and posture [Argyle, 1988, Burgoon et al., 1984, Burgoon and Le Poire, 1999, Richmond et al., 2008], perceivers can determine to what extent another

Dimensions	High Value Traits	Low Value Traits
Extraversion	Sociable, Friendly, Talkative, Fun-loving	Introverted, Reserved, Inhibited, Quiet
Agreeableness	Courteous, Forgiving, Sympathetic	Critical, Rude, Harsh, Callous
Conscientiousness	Reliable, Careful, Well-organized, Self-disciplined	Negligent, Disorganized, Undependable
Neuroticism	Nervous, Insecure, Worrying, High-strung	Calm, Relaxed, Secure, Hardy
Openness	Creative, Curious, Complex	Conventional, Narrow interests, Uncreative

Table 2.1: The components of the Big 5 model and high/low value traits for each dimension.

individual is generally extroverted or introverted [Nalini and Skowronski, 2008, p. 122].

Those impressions have been shown to be accurate when observers were exposed to relatively short “thin slices” (4-10 min) of ongoing streams of individuals’ behavior [Ambady and Rosenthal, 1992]. Levesque and Kenny, for example, conducted a study [Levesque and Kenny, 1993] where groups of four strangers rated each other on the Big 5. Then, the strangers met in pairs and were videotaped talking to each other. Later, judges watched the extensive videotapes and rated each subject’s extroversion, based on the amount of time he or she spent talking, the number of arm movements and other factors. As a result, the strangers’ first-impression ratings of extroversion strongly correlated with people’s rated levels of extroversion as seen on the videotape.

Furthermore, perceivers update their impressions as the sample of available evidence increases during the course of acquaintanceship. Shifts in impressions, though not inevitable, might happen in developing relationships. However, impressions of extraversion seem to require an abundance of evidence to move substantially in either direction, therefore once they are formed they benefit of great stability and are easy to keep up [Kammrath et al., 2007].

Whereas personality represents the unique characteristics of an individual [Miller et al., 2007, p. 26] and arise out of more indirect and long-term factors [Moffat, 1997], interpersonal attitudes of a person toward another are temporally inconsistent and subject to a greater degree of variation, they can vary in different contexts and have shorter duration.

Argyle's Model of Interpersonal Attitudes

Interpersonal attitudes are essentially an individual's conscious or unconscious evaluation of how they feel about and relate to another person [Argyle, 1988, p. 85]. Attitudes towards other people are rather similar to emotions and may involve exactly the same signals. For example, being angry is an emotion, being friendly or hostile towards someone is an interpersonal attitude. Therefore, compared to personality, attitudes are temporally inconsistent and subject to a greater degree of variation. They can be expressed by body language in a number of ways. Argyle identifies two fundamental dimensions that can account for a great variety of nonverbal behavior (i.e. proximity, gaze, posture, touch, facial expression, etc . . .) **affiliation** (ranging from friendly to hostile) and **status** (from dominant to submissive) [Argyle, 1988, p. 86].

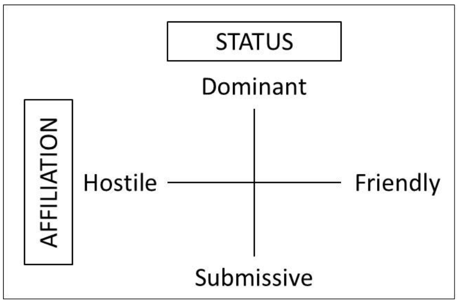

Figure 2.1: The two dimensions of attitudes toward others described by Argyle's status and affiliation model.

Affiliation can be broadly characterized as liking or wanting a close relationship, whereas status is the social superiority (dominance) or inferiority (submission) of one person relative to another (as depicted in figure 2.1). Attitudes and their expression can depend both on the general disposition of the person and their relationship to the other person, for example status depends on whether they are generally confident of their status and whether they feel superior to the person they are with. On the other hand, controlling for affiliation, i.e being friendly, helps people to establish a relationship with another person or show politeness during first encounters [Argyle, 1988, p. 86].

There are managed aspects of interpersonal attitudes, in particular for the affiliation dimension, that someone might want to control to make sure that the exhibited behavior during a greeting encounter is consistent with the desired impression of attitude aimed at managing toward others.

In general, first impressions of either personality or interpersonal attitudes do not happen in the void. In the specific context of greeting encounters, there exist precise social norms that regulate the interaction between participants [Kendon, 1990], and the nonverbal behavior is exhibited in relation to the physical space that separates them [Hall, 1966]. Considering an encounter where a person is getting closer to another, the decreasing interpersonal distance becomes a trigger on top of which nonverbal behavior occurs at specific points and, as a consequence, first impressions are shaped based on the observed behavior [Argyle, 1988].

2.6 Interpersonal Distance and Human Greetings

Hall's Proxemics Theory

Interpersonal distance can be seen as an invisible set of bubbles surrounding us [Miller et al., 2007]. According to [Hall, 1966], humans tend to exhibit different behaviors towards each other in accordance of four levels of "closeness" depending on the type of interpersonal relationship we are involved in with our interaction partner. He labeled these four levels as intimate, personal, social and public zones.

- *Intimate zone* (people are 0 to 0.45 meters apart): This zone tends to be reserved for people within intimate relationships (e.g. close family members or lovers). Interactions in this space tend to be physical and vocalizations are minimal.
- *Personal zone* (from 0.45 to 1.2 meters): This zone is mainly used during conversations with close friends and interactions with relatives. It is possible to touch the other by reaching in this zone.
- *Social zone* (from 1.2 to 3.6 meters): This is often the distance we allow acquaintances to stay in. Interactors use a higher level of gaze, body movements are visible.
- *Public zone* (beyond 3.6 meters): This is the outer zone and sustained interaction mainly happens at this level in the context of presentations and public figures. This distance also marks the boundaries between us and people we wish to stay away from and not interact with.

These distances describe the general patterns of interactions in western cultures [Miller et al., 2007]. Nonverbal behavior in some circumstances, such as in greeting encounters, might not appear in a discrete fashion associated within each proxemics zone, but it is likely to be part of an overall continuous process that spans across multiple zones. In fact how far one "goes out of one's way" to meet another appears to have a precise communicational significance [Kendon, 1990, p. 179] as thoroughly described in Kendon's salutation model of human greetings.

Kendon's Greeting Model

Kendon observed and videotaped six different social situations where greetings would take place, and he noted down greeting behaviors and stages which commonly occurred, in particular for dyadic greetings between adults. He refers to "greeting" as the unit of social interaction often observed when people come into one another's presence [Kendon, 1990, p. 153].

With his observations Kendon discovered that two individuals that wish to greet each other adhere to a structured interaction comprising the following phases: sighting, distant salutation, approach, close salutation and initiation of conversation. These phases, as described below, come with an implicit relation to proxemics. In fact specific behaviors are displayed by greeters at inexact but predictable distances.

The following is a description of Kendon's observations in each phase. We will refer to a specific instance of a greeting between two persons p and w, where p is walking towards w and w is waiting standing still.

Sighting, recognition and decision to greet. Before any greeting can begin, the participants must sight each other, and in doing so they must identify the other as someone they wish to greet. This usually happens by the means of a short exchange of glances and w orienting the body towards p as invitation cue to begin the greeting process [Kendon, 1990, p. 167]. From this stage on, behaviors observed depend on the relationship between greeters, thus for example strangers in zero-acquaintance situations might go through all of the following phases, whereas acquaintances or friends upon recognizing each other could proceed to the final phase of the approach not necessarily following all the rituals of common behavior described as follows [Kendon, 1990, p. 203].

Distant salutation. Immediately following the sighting the approach begins with a distant salutation at approximately 8 meters. This phase establishes that the two greeters have seen one another and that they are now ratified in a greeting process. It occurs at such a distance that the greeters, if they continue to interact, move closer to one another afterwards and commit to engage in interaction. [Kendon, 1990, p. 172]. A number of different behaviors may be observed, in particular w might *smile* and this may or may not continue to the next phase. A w's *wave* is common, it may vary according to the distance, but usually is in the form of hand raised and palm open and oriented towards p [Kendon, 1990, p. 175]. A variety of *head behaviors* also occur:

- The *head toss*: the head is tilted back rapidly, and then brought forward again. Sometimes is coupled with an almost imperceptible eyebrow raise [Kendon, 1990, p. 174];
- The *head nod*: the head is lowered and then at once raised [Kendon, 1990, p. 175];
- The *head lower*: the head is lowered or tilted forward, and it is distinguished from the nod only in its duration [Kendon, 1990, p. 174];
- The *head dip*: by one of the participants often follows one of the above behaviors, where a person lowers their head to mark a shift of attention, for example, when p is moving into the next phase of the greeting [Kendon, 1990, p. 177].

Head toss, lower and nod are also linked to p's recognition by w in terms of relationship status, thus in greeting situations where p and w are strangers they might not appear, whereas a head nod indicating immediate recognition might occur between a pair of acquaintances or friends [Kendon, 1990, p. 175].

Approach. In most of the greetings that include a close salutation we shall observe an approach [Kendon, 1990, p. 179]. In proxemics terms, greeters move towards one of the zones most appropriate to their interaction (e.g., public, personal or intimate). An interesting *pattern of gazes* fills the distant and close salutation phases during the approach. People do not look at each other continuously as they approach one another [Kendon, 1990, p. 180]. Although within half second immediately preceding the start of the close salutation almost all the greeters are "looking at" one another, during the previous three seconds the proportion of time spent looking is sharply lower. In general , we could say that as the approach begins the greeters tend to look at each other as distant salutation occurs. This is often followed by a head dip and thereafter, some looking is likely to occur but less and less as the approach continues until, when the individuals are within a few meters of one another they look comparatively little with a sharp "cut-off" right before the close salutation. Though looking away increases sharply in most cases just before close salutation, thereafter looking at increases even more sharply. Almost everyone is looking directly at the other as the close salutation begins. Furthermore, looking away tends to happen with a change of facial orientation (or gaze) [Kendon, 1990, p. 180-182].

Some behavior intended for regulation of intimacy have also been observed in this phase. A "*body cross*", that is when either p or w brings one or both arms in front of him, “crossing” the upper part of the body [Kendon, 1990, p. 185]. *Grooming* is also likely to be observed in this stage of the greeting, people may adjust hair or clothing [Kendon, 1990, p. 185].

Final approach. This phase immediately precedes the start of the close salutation [Kendon, 1990, p. 188]. People commonly orient their palms towards those they are greeting, in our example w, in the “palm presentation” gesture [Kendon, 1990, p. 191]. If a person is not yet smiling during the approach, he or she will typically smile during the final approach [Kendon, 1990, p. 188]. Smile is commonly associated with the distance salutation, however, thereafter the smile either rapidly fades in intensity, or even disappears altogether, until participants are close to one another and once again are looking at one another. Very often, however, a definite smile could not be distinguishable again until the final approach phase [Kendon, 1990, p. 189]. Both greeters are likely to hold their head in a way that is distinct from the way they have been holding it in earlier phases. Kendon describes the different head position in this phase as headset [Kendon, 1990, p. 189].

Close salutation. At approximately 3 meters away greeters engage in a close salutation. This greeting phase marks the end of the approach but also establishes the relationship that the greeters have, thus we can expect that aspects such as friendliness and identity will be signaled at this stage [Kendon, 1990, p. 203]. Up till now participants have substantially reduced the distance between them [Kendon, 1990, p. 191] and are capable of being in physical contact with each other, however different forms of salutations may occur:

- *Non-contact close salutation*. In this case, participants halt facing one another, but apart from exchanging verbal greetings they simply sustain eye contact, the headset and facial display they had assumed during the approach. Also body postural change is observed, such as a posture shift or a lean forward. [Kendon, 1990, p. 193].

- *Handshakes*. They vary in length and intensity and are influenced by participants’ gender and formality of the occasion [Kendon, 1990, p. 195].

- *Embraces*. They highly depend on the relationship between the two greeters [Kendon, 1990, p. 197].

- *Other forms*. These vary between cultures. Some examples include bowing, cheek to cheek kisses, etc . . . [Kendon, 1990, p. 199].

Encounter: The Initiation of Conversation. Finally this is the stage, named by Kendon the *"how are you phase"*, where people exchange information about one another and the greeting process transitions to a face-to-face conversation [Kendon, 1990, p. 203]. As discussed in proxemics, once the close greeting has been performed, participants tend to increase or decrease the space between them in a way that matches the nature of their interaction, where they may step back from the intimate zone to the personal or public zone.

In conclusion, greetings serve as precursor leading to interaction [Kendon, 1990, p. 154], and how many stages of the greeting program are enacted will depend upon the kind of relationship that exists between the greeters. For a given kind of relationship there are situations where all the steps of the program are appropriate, but others in which only some of them will be necessary. In situations where close interaction is to follow, for example, it appears that upon an initial encounter (i.e. the first time two individuals encounter one another within the context of a given occasion) both a distance and a close exchange salutation will occur. Whereas upon subsequent encounters only a distant salutation is necessary [Kendon, 1990, p.203].

2.7 Conclusions

In this chapter we described the theoretical models adopted for our approach with more details. The Kendon's greeting model and Hall's proxemics theory serves as basis on which all the others will be applied in the design of our empirical studies described in chapter 4.

The idea is to manipulate the agent's nonverbal immediacy (i.e. smile, gaze and proxemic behavior) during the phases of the greeting encounter with the user, and understand what are the impressions of extraversion (according to the Big 5 model) and friendliness (according to Argyle's status and affiliation model) that users form. Subsequently, we study whether those impressions have an impact on relational decisions about encountering again the agent as suggested by the predicted outcome value theory.

Prior to describing our theoretical framework, the next chapter will cover related works in the field of embodied conversational agents and relational agents.

"Science is a bit like sports: it cannot deny certain results, provided that, obviously, they are judged by independent referees."

Piero Angela (1928 – present)

3

Related Agent Work

In this chapter we review prior work on relational agents and embodied conversational agents involving the background psychological and sociological concepts introduced earlier. These works represent the foundations for the theoretical framework that we will describe in the next chapter and they are all focusing on agents' nonverbal communicative behavior, with particular emphasis on smiling, gazing and proxemic immediacy cues.

We first describe some work dealing with agents' impression management and formation in Section 3.1, then in Section 3.2 we briefly show relational agents work that focuses on the long term impact, in terms of relational goals, of the agent's nonverbal behavior exhibited. In Section 3.3 we show ECA systems aimed at expressing personality traits and interpersonal attitudes (separately from each other). In Section 3.4 we analyze the work dealing with user's attributes, with emphasis on user's personality, and their impact on the agent's evaluation and the user's preferences for a specific agent nonverbal behavior. Section 3.5 reviews user's perception mechanisms based on interpersonal distance and proximity inputs. Finally, Section 3.6 shows some deployments of ECAs and RAs in public spaces with emphasis on museum agents.

3.1 Impression Formation and Management

[Maat and Heylen, 2009, Maat et al., 2010] showed how a realization of a simple communicative function (managing the interaction) could influence users' impressions of an agent. They focused on impressions of personality (agreeableness), emotion and social attitudes through different turn-taking strategies in human face-to-face conversations applied to their virtual agents in order to create different impressions of them. [Fukayama et al., 2002] proposed and evaluated a gaze movement model that enabled an embodied interface agent to convey different impression to users. They used an "eyes-only" agent on a black background and the impressions they focused on were affiliation (friendliness, warmth) and status (dominance, assurance). Similarly, [Takashima et al., 2008] evaluated the effects of different eye blinking rates of virtual agents on the viewers subjective impressions of friendliness, nervousness and intelligence.

The work of Heylen et al. emphasizes the "side-effect" of different nonverbal choices in the realization of a communicative function (i.e. turn taking), whereas our purpose is to intentionally manipulate specific agents' immediacy cues (smile, gaze and proximity) and see how users interpret them. The interest is on the impressions they form of personality/affiliation but also keeping an eye on extra types of judgments that could arise. As opposed to [Fukayama et al., 2002], we are using full body virtual agents to exhibit our nonverbal behavior, which is not narrowed down to specific behaviors such as eyes-only gaze [Fukayama et al., 2002] or eyes blinking [Takashima et al., 2008].

3.2 Long-term Impact of Nonverbal Immediacy

There are a significant number of studies on relational agents, in different application domains, investigating the importance of an agent's nonverbal behavior when having a simulated face-to-face conversation with the user. The underlying importance of these studies concerns the impact that the behavior exhibited might have on the effectiveness of the agents, in particular on the agents ability to establish and carry on a relationship with their users, therefore the ability to work together with them. We survey here a few key findings.

[Kang et al., 2012] presented a study of human nonverbal behavior during intimate self-disclosure in interviews and applied their findings to the design of a virtual counselor. They argued that expressing different intimacy levels accompanied by proper nonverbal behavior is fundamental for a virtual counselor to induce reciprocal disclosure in the human clients and therefore enabling them to like their virtual counselors more and create better rapport with them.

[Wang and Gratch, 2009] investigated the effectiveness of virtual human nonverbal immediacy behavior in creating rapport with a human learner and promoting learning. Although they were not able to demonstrate a direct link between the two, their results suggested that creating rapport relates to higher self-efficacy, which is related to better learning results.

[Schulman and Bickmore, 2012] presented a nonverbal behavior generation system for a relational counselor agent. Their computational model was based on an observational study of multiple dyadic conversations between a health counselor and her clients.

Our work has foundations in the nonverbal immediacy cues exhibited by a relational agent as in the works described, but we narrowed down the attention to user-agent interactions in the context of first time greeting encounters. In doing so, we are particularly interested in the impact that some impressions of personality and attitudes the user might form of the agents in these very brief encounters can have in the longer-term, when the user decides whether to interact again with the agent or not.

Most recently, [Bergmann et al., 2012] investigated how appearance and nonverbal immediacy affect the perceived warmth and competence of virtual agents over time. Their goal was to study how warmth and competence ratings changed from a first impression after a few seconds to a second impression after a longer period of human-agent interaction, depending on manipulations of the virtual agent's appearance (robotic vs. human like) and nonverbal behavior (focusing on gestural behavior). This work complements our research, but the main differences consist of (a) the manipulations adopted, in our case nonverbal immediacy cues whereas they focused on gestures; (b) the dimensions along which first impressions were measured, in our case are personality and attitude; (c) their application domain was a task-related monologue that required visual descriptions, whereas ours is a virtual greeting encounter with an agent and (d) their main goal was studying impressions formed when interpreting nonverbal cues, as in our work, but they did not examine the impact on users' relational decisions.

3.3 Expression of Personality and Interpersonal Attitudes

There has been considerable previous work developing expressive virtual characters capable of reflecting a personality consistent with the verbal and nonverbal cues exhibited. Neff et al. exploited the extraversion [Neff et al., 2010] and neuroticism [Neff et al., 2011] traits of the Big 5 model [Goldberg and Saucier, 1998] in multimodal virtual characters evaluating the effects of verbal and nonverbal behavior in personality perception studies. Similarly, [Doce et al., 2010] presented a model of personality, based on the Big 5, aimed at creating distinct traits that in turn can influence an agent's cognitive and behavioral processes. [Sevin et al., 2010] proposed a real-time back-channel selection algorithm for choosing the type and frequency of back-channels to be displayed according to the personality of the virtual character used (shown in Fig. 3.1a).

How should I look like when I **deny** something with a submissive attitude?

Now, can you help our character to deny something with a submissive attitude?

Select the appropriate parameters!

Select your parameters

Facial expression
○ No expression ○ Smile ● Frown

Gesture
○ No ○ Head movement ○ Gesture ● Both

Amplitude of gesture
○ Small ○ Medium ● Large

Power of gesture
● Small ○ Medium ○ Strong

Head position
○ Straight ○ Down ○ Up ● Tilt

Gaze
● Gaze at ○ Gaze away

Are you satisfied with the result?
Not at all ○ ○ ○ ○ ○ ○ ○ Great !
Average

(a) Expression of Personality (b) Expression of Attitudes

Figure 3.1: Expression of personality and interpersonal attitudes in virtual agents. Figure (a) depicts the listener agent used by [Sevin et al., 2010] to evaluate their algorithm for selecting back-channel behavior according to the agent's personality. Figure (b) shows the interface used by [Ravenet et al., 2013] to collect the corpus of a virtual agent's nonverbal behavior conveying different interpersonal attitudes.

Regarding interpersonal attitudes, [Ballin et al., 2004] concentrated on a general framework based on Argyle's status and affiliation model [Argyle, 1988] for animating nonverbal behavior of virtual characters in improvisational visual media production and expressing interpersonal attitudes towards one another. [Lee and Marsella, 2011] proposed an analysis framework of nonverbal behavior for modeling side participants and bystanders. They based their analysis on Argyle's status and affiliation model and considered agents' interpersonal relationships, communicative acts and conversational roles. [Ravenet et al., 2013] adopted a user perceptive approach to modeling the nonverbal behavior that an agent should exhibit to convey interpersonal attitudes (both status and affiliation dimensions of Argyle's model) during face-to-face conversations with the users. They first collected a corpus using an online platform where users directly configured the behavior that they would expect from an agent for conveying different attitudes (see Fig. 3.1b). Then they developed a Bayesian network based on the data gathered to create a computational model for autonomous multimodal behavior generation.

These works dealt with either incorporating personality traits [Neff et al., 2010, Neff et al., 2011, Doce et al., 2010, Sevin et al., 2010] or interpersonal attitudes [Ballin et al., 2004, Lee and Marsella, 2011, Ravenet et al., 2013] separately. The virtual agents were mainly designed for face-to-face interactions or interactive drama. These works demonstrated a practical usage of human social theories for modeling personality traits and interpersonal attitudes, respectively, with the Big 5 and Argyle's status and affiliation models. We also adopted these models in order to share a common theoretical background and facilitate the comparison between previous agent work and ours in terms of agents' evaluation. However, our work focuses on user's judgments of nonverbal behavior when both personality (extraversion) and attitude (affiliation) are expressed at the same time. Furthermore, in our framework the agents are exclusively exhibiting nonverbal behavior in the formative moments of the first virtual encounter between the user and agent.

3.4 Impact of User's Attributes

[Krämer et al., 2010] demonstrated how inter-individual differences can lead to preferences for specific nonverbal behavior exhibited by interface agents. They investigated whether an agent's nonverbal behavior had different effects on different users in a series of evaluation studies deploying life-sized ECAs interacting with the users. The agents' nonverbal behavior tested was eyebrow movements, self-touching gesturing, frequency of smiling and different back-channel feedback styles. The effects measured were users' feelings during the interaction with the agent (e.g. "attentive", "relaxed", etc. . .), person perception of the agent (e.g. "self-confident", "cold", etc. . .), and general evaluation of the interaction with the agent (e.g. user's enjoyment). They studied the moderating effects of participant's gender, age and computer literacy and concluded that all of the three attributes had influence, among their findings, on the feelings of the participants (e.g. women were more nervous than men when interacting with the agents varying the frequency of smiling behavior) and on the agent's evaluation (e.g. for computer literacy, experts experienced self-touching agents as more strained compared to novices). In conclusion, their important suggestion was to tailor an agent's nonverbal behavior to different user groups.

We also aim to take into account user's attributes in the evaluation studies included in our theoretical framework, however, while they focused on user's gender, age and computer literacy, our focus is on the user's own personality and the impact that it might have on (a) the overall users' preferences for a specific agent attitude and personality and (b) on users evaluation of greeting agents' nonverbal behavior.

Impact of User's Personality on Agent Preference

Studies in human-robot and human-agent interaction that targeted the benefits of a match between agent and user personality have yielded inconsistent results. [Isbister and Nass, 2000] found that people tended to prefer a (static) virtual character figure whose personality was complementary to their own, while [Nass et al., 1995] (computer interface) and [Tapus et al., 2008] (robot) showed that people preferred to interact with personalities similar to their own. [Bickmore and Cassell, 2001] showed that an agent using small talk when interacting with users increased trust in it for those extraverted, but for the introverted it had no effect.

Given that we expect users to react to social cues from computers or robots in ways that are similar to how they would react to the same cues from a person [Nass et al., 1995], we want to investigate whether users prefer virtual agents with personalities similar to or complementary to their own personality. Our focus is on embodied conversational agents (relational in particular) and on the specific context of first impressions where our agents exhibit nonverbal immediacy cues to engage people and express a specific personality trait (extraversion) and a clear interpersonal attitude (friendliness) with their nonverbal behavior.

Impact of User's Personality on Agent Evaluation

According to [Pütten et al., 2010], users' personality influences their subjective feeling and actual behavior after the interaction with a virtual agent, as well as how they evaluate the agent. The effects of an agent's behavior also depends on the personality of the user, in particular people with high values in agreeableness and extraversion (among other findings) judged agents more positively compared to people with high values in shyness. Kang and colleagues suggested that users' personality traits crucially affect their perceptions of virtual agents. They explored how users' shyness [Kang et al., 2008b] and Big 5 personality traits [Kang et al., 2008a] are associated with their feelings of rapport when they interacted with different versions of virtual agents capable of exhibiting nonverbal feedback. [Kang et al., 2008b] found that more anxious people (high in social anxiety, i.e. shyness) felt less rapport, while feeling more embarrassment, when they interacted with a *non-contingent* agent. On the other hand, according to [Kang et al., 2008a], more agreeable people felt strong rapport when interacting with a rapport agent embodying agreeable features (i.e. nonverbal *contingent* feedback while listening).

As opposed to the typology of studies investigating the benefits of matching-up user and agent personality (e.g. [Isbister and Nass, 2000]), we aim to understand the role of a user's personality when interacting with a virtual agent, similar to [Kang et al., 2008b, Kang et al., 2008a]. However, in our context we are interested in the possible blending effect that user personality may have on snap judgments of personality/affiliation after observations of solely body language in the very first moments of interaction.

3.5 Interpersonal Distance and User Perception

In this section we review some virtual and robotic agents work dealing with user detection and perception. In particular these agents used interpersonal distance as input cue to understand when the user was willing to initiate an interaction or, in some cases, to react towards approaching users by exhibiting different nonverbal behavior in line with Kendon and Hall's theories [Hall, 1966, Kendon, 1990].

In *BodyChat* [Vilhjálmsson and Cassell, 1998] users could chat with each other by using text while their avatars in a shared online 3D environment automatically animated behaviors to accomplish important communicative functions such as attention and salutations (among the others). In particular, an user's availability to chat was communicated via his/her own avatar by exhibiting an initial exchange of glances and performing distant and close salutations towards another approaching avatar (as suggested by Kendon's greeting model) that showed interest for chatting (see Fig. 3.2b).

REA (Real Estate Agent) [Cassell et al., 2000a] was a relational agent who acted as a real estate salesperson that could take the user on a virtual tour of some properties and point out and discuss their salient features. The agent, shown in Fig. 3.2c, was capable of both multimodal input understanding and output generation. She acknowledged the user's presence through a vision system that was tracking both the user's head and hand position in space during conversation. A microphone was used for capturing speech input. Greetings (and farewell) with the user were generated, but only in response to the user's verbal greeting and farewell. A fully articulated graphical life-sized body allowed the agent to exhibit eye gaze, body posture, hand gestures, and facial displays to organize and regulate her interaction.

MACK [Cassell et al., 2002] was a life-sized embodied conversational agent who could answer questions about and give directions to the MIT Media Lab's various research groups, projects, and people. The agent was deployed in the form of a kiosk, it was capable of detecting the users' presence, though only when they were in close proximity of the agent's installation by way of a pressure-sensing chair mat.

Tinker [Bickmore et al., 2013], the human-sized relational museum guide installed at the at the Boston Museum of Science (described with more detail in Section 4.4) originally used a motion sensor to determine if visitors were in the surrounding general area so that it could beckon them over to talk and begin conversation initiation behaviors.

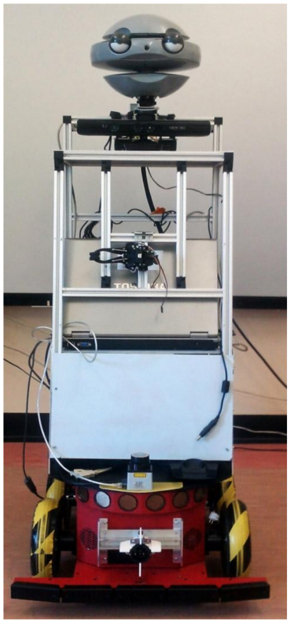

(a) Robotic Companion

(b) BodyChat

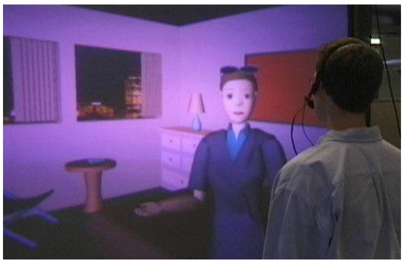

(c) REA

Figure 3.2: Examples of virtual and robotic agents operating on user's detection and intepersonal distance. Figure (a) shows the robotic RA companion deployed by [Deshmukh et al., 2010] in their lab facilities. The robot used the Micorsoft Kinect technology to detect researchers willing to engage in interaction. Figure (b) shows another user's avatar in the *BodyChat* system while glancing at the user as part of the greeting process. Figure (c) shows the virtual real estate agent *REA* while interacting with a user detected by her vision and tracking systems.

[Ring et al., 2013] developed a relational agent for isolated older adults to explore techniques of providing automated social support over extended periods of time. This agent was meant to be used within a user's home via a touchscreen computer. They developed two initial exploration versions named *passive* and *proactive*. In the passive version, conversations with the agent had to be initiated by the user by touching an option on the touchscreen, whereas in the proactive version, the agent could detect when the users walked by the system via a motion sensor and attempted to initiate a conversation by verbally greeting them. Results from an exploratory pilot study indicated that when the agent proactively drawn elders into interactions, it was more effective at addressing loneliness than when the agent passively relied upon elders to initiate interactions.

[Deshmukh et al., 2010] worked on a companion robotic relational agent capable of sharing a lab with a small group of researchers and performing tasks such as carrying the phone and printed material to users or giving out reminders for important events (e.g. meetings). The robot, depicted in Fig. 3.2a, had user's face tracking and motion direction detection abilities. These perceptions enabled it to engage in spontaneous interaction with the users by approaching them in a proper manner (i.e. asses whether the interaction can be initiated even when not requested) or be ready to start an interaction if the user showed willingness to engage with the robot. In the latter scenario, the agent after detecting the user's presence, it responded engaging in an interaction only when a certain social area (as defined by Hall's theory) was reached by the user.

[Heenan et al., 2013] adopted Kendon's greeting model and Hall's proxemics theory to provide a humanoid (Nao) robotic agent with social awareness towards approaching users. The robot exhibited some of the nonverbal behavior described in Kendon's model to accomplish the greeting communicative functions and increase its social believability. The physical constraints of the robot used prevented it from realizing facial expressions and eye gaze behavior, however proxemics, gaze (whole head) and simple gesturing (e.g. hand wave) were exhibited during the greeting process with the user.

All the works presented except for [Vilhjálmsson and Cassell, 1998, Cassell et al., 2002, Heenan et al., 2013] were relational agents. Some agents had a virtual embodiment [Vilhjálmsson and Cassell, 1998, Cassell et al., 2000a, Cassell et al., 2002, Bickmore et al., 2013, Ring et al., 2013] and others were robotic agents [Heenan et al., 2013, Deshmukh et al., 2010]. Independently from the agent embodiment, they all suggest that considering interpersonal distance between user and agent and, in particular, paying attention to the nonverbal behavior exhibited by the agent during the initial moments of interaction (e.g. the greeting process) are fundamental aspects to consider when building RAs (and ECAS in general).

The two works not concerning the relational agents domain were included in this review since they both represent direct applications of Kendon's greeting model as in our approach. However, [Vilhjálmsson and Cassell, 1998] dealt with avatar's behavior automation in the computer mediated communication domain and in [Heenan et al., 2013] they dealt with robots. Both applications adopted Kendon's model to aid the initiation of interaction with the user and perform salutations, we also apply this theory in a similar fashion but we are interested in manipulating the subtle cues exhibited during the greeting process.

All the relational agents that we discussed were able to detect the user presence in some way, mainly with motion detection sensors. However, none of them considered the opportunity to begin the interaction with the user from distance and exhibit different nonverbal behavior according to the decreasing interpersonal distance as the user is getting closer to the agent.

MACK [Cassell et al., 2002] represents the only non-relational agent work in our review that had a certain degree of user awareness with a pressure sensor on a chair, however it did not employ any user perception technique based on interpersonal distance (i.e. detecting approaching users).

3.6 Public Space Deployments

A specific goal of our theoretical framework, that we will introduce later, is to test the effects of managing first impressions of a relational agent in a real application setting. The work reviewed in this section refers to deployment of both virtual and robotic conversational agents in public space exhibitions, with particular emphasis on interactive museum installations.

[Gockley et al., 2005] presented *Valerie*, a robot receptionist installed in the entrance way to Newell-Simon Hall at Carnegie Mellon University, in USA. Valerie was able to give directions to visitors and lookup the weather forecast while also exhibiting a compelling personality and character to encourage multiple visits over extended periods of time.

[Kopp et al., 2005] installed *Max* in the Heinz Nixdorf Museums Forum (HNF), in Paderborn (Germany). Max was projected on a life-size screen, standing face-to-face with visitors and acting as a museum guide. It was able to engage with visitors in natural face-to-face conversations to provide information about the museum, the exhibition, or other topics of interest. Its responses to visitors included a German voice accompanied by appropriate nonverbal behaviors like facial expressions, gaze, or locomotion.

[Traum et al., 2012] developed the *Ada and Grace* exhibit, installed near Tinker in the Boston Museum of Science. Ada and Grace are twin conversational agents able to interact with each other and are designed to engage visitors using unrestricted speech input from a microphone and to discuss various science topics. Figure 3.3 shows the agents' deployment at the Museum of Science.

Figure 3.3: An example of virtual agents deployed in a public space. The image shows the *Ada and Grace* exhibit installed at the Boston Museum of Science.

[Lane et al., 2011] discussed *Coach Mike*, a 3D cartoon-style pedagogical agent at the Boston Museum of Science that sought to help visitors at Robot Park, an interactive exhibit for computer programming. Mike was designed to be approachable, supportive, and understanding in order to improve the experience of museum visitors in that area.

Finally, [Al Moubayed et al., 2012] deployed *Furhat* at the London Science Museum as part of a robot festival. Furhat is a three dimensional back-projected human like robot head that utilizes a computer animated face to carry multimodal multiparty interactions with users. In the museum deployment, in particular, it had the task of asking the visitors about their beliefs of the future of robots, with the possibility of talking to two visitors simultaneously shifting attention between them.

The work described incorporates features directly related to the theoretical framework that we will propose in the next chapter. The goal was to build relationships between the agents and their users (except for [Traum et al., 2012, Lane et al., 2011, Al Moubayed et al., 2012]), to interact with them using multiple modalities, and to engage the visitors during the current visit or encourage them to come for multiple visits over extended periods of time. However, none of them incorporates a model to manage first impressions through nonverbal behavior when visitors approach them for the first time.

"All truths are easy to understand once they are discovered,
the point is to discover them."
Galileo Galilei (1564 – 1642)

4

Theoretical Framework

4.1 The Idea: A Combination of Theories

So far, the importance of observed nonverbal behavior in the context of impression formation of someone's personality and interpersonal attitude in greeting encounters has been examined, and it has been explained how those impressions link to relational decisions (in chapter 2). The related works described in chapter 3 considered several of these aspects in the implementation (see Sections 3.5 and 3.6) and evaluation (see Sections 3.1, 3.2 and 3.4) of relational agents and embodied conversational agents in general, but a novel unified theory has not been provided yet.

This thesis proposes a framework that consolidates all the theories that have been discussed so far for building more effective relational agents. Figure 4.1 depicts a conceptual view of the framework. Kendon's greeting model and Hall's proxemics theory combined together constitute the infrastructure. A relational interface agent should encourage and invite interaction by exhibiting the appropriate nonverbal communicative behavior early in the very first steps of the greeting process according to the phases observed by Kendon, and do this while observing Hall proxemics zones and the interpersonal distance between the user and the agent.

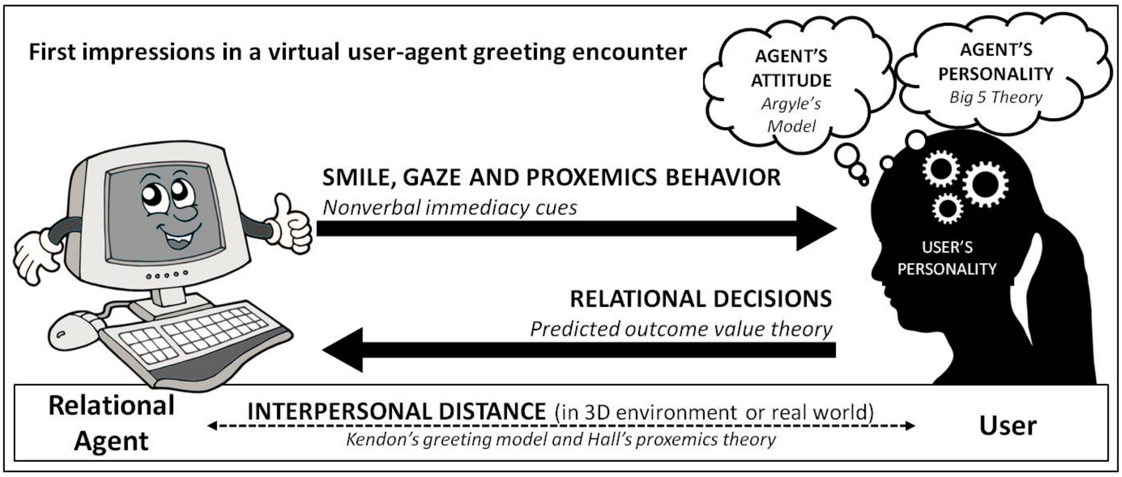

Figure 4.1: A conceptual view of the combination of theories involved in our theoretical framework grounded in Kendon's and Hall's theories. As the interpersonal distance between user and agent decreases (either in a 3D environment or the real world), the nonverbal immediacy cues of smile, gaze and proxemics exhibited by the agent is interpreted by the user and, depending also on his/her own personality, contribute to shaping impressions of the agent's personality and attitude. As outcome, these impressions impact users' relational decisions about further interactions with the agent.

Either in a virtual 3D environment or a real life deployment, the interpersonal distance of the user's avatar, in the former scenario, or the physical position in the real world in the latter, represents a valuable triggering mechanism for agent's nonverbal reactions. On top of this, subtle manipulations of agent's nonverbal immediacy cues of smile, gaze and proxemics, in addition to serve basic greeting communicative functions, they also account for impressions of agent's extraversion (personality trait according to the big 5 model) and affiliation (interpersonal attitude according to Argyle model).

There are at least three main issues that need to be exploited when it comes to interpreting those immediacy cues when exhibited in concert. First, this interpretation is not always obvious, since each behavior can account for either impressions of personality or attitude. Secondly, the interpretations can be also affected by the actual deployment of the agent. In case of 3D virtual environments, for example, the camera perspective used to render the scene can alter the way nonverbal behavior is interpreted. This perspective can be either in first person view (1P) or third person view (3P). The difference between the two, in videogames environments (but it can be easily generalized to all 3D environments), is that first person refers to a graphical perspective rendered from the viewpoint of the player's controlled character (avatar), whereas third person refers to a graphical perspective rendered from a view that is some distance away (usually behind and slightly above) from the avatar. Thirdly, does the interpretation of nonverbal behavior change according to the user's own personality?

A relational agent might be able to properly carry on a simulated face-to-face conversation and make use of its relational skills, but the formative moments of the very first interaction with the user represent a *reading key* for what to expect from it later, in a way those moments give the users a taste of how further interactions with the agent would be. The impressions that users form may be predictive of future relational decisions, in terms of how likely they would be to interact with the agent again and how often they would interact with it, according to the predicted outcome value theory. However, the agent's extraversion level combined with the affiliation expressed towards the user constitute a great deal of information that combined may lead to unpredictable outcomes. Considering the different ways personality and attitudes endure over time, we aim to investigate how the two combine together and whether one of the two has greater impact on users relational decisions.

Our framework addresses the complexity of the issues presented by separating our research questions across three user evaluation studies:

1. **Nonverbal Behavior Interpretation Study.** We evaluate users' impressions of a greeting agent's extraversion and affiliation as a result of interpreting nonverbal immediacy cues of smile, gaze and proxemics in a first virtual encounter (in Section 4.2) [Cafaro et al., 2012].
2. **Nonverbal Behavior Impact Study.** We investigate whether the first impressions of extraversion and affiliation that users form of a relational agent have an impact on subsequent relational behavior (in Section 4.3) [Cafaro et al., 2013].
3. **Managing First Impressions in a Public Space Study.** We move out of the virtual world to test the effectiveness of managing first impressions for a relational agent installed in a real public setting, with thousands of people approaching it, therefore working with physical interpersonal distance between the user and the agent (in Section 4.4).

4.2 Nonverbal Behavior Interpretation Study

In order to evaluate users' impressions of a greeting agent's extraversion and affiliation in a first encounter we conducted this empirical study in which subjects, represented by an on-screen avatar, approached a series of greeting agents in a 3D virtual museum entrance displayed on a 19″ LCD monitor. The agents exclusively exhibited a set of nonverbal immediacy cues that were systematically manipulated during approaches of 12.5 seconds each (the duration was chosen after a prior validation study described in the appendix A.1).

The main research questions were the following:

1. What is the role of smile, gaze and proxemics behavior when managing impressions of extraversion and affiliation?
2. How do those cues combine in user interpretations?
3. Does the interpretation of nonverbal behavior change according to users' own personality?

Furthermore, we had two additional exploratory questions: (1) In addition to personality and interpersonal attitudes, what kind of judgments are people forming when observing and interacting with RAs showing only specific nonverbal cues? (2) When moving from a first to third person camera perspective, do users interpret observed nonverbal behavior in the same way? For this reason the study was split in two trials, 1P and 3P, differing only in the camera perspective used.

4.2.1 Hypotheses

Our hypotheses, for both trials, were the following:

- **H1**: The amount of extraversion that subjects attribute to a greeting agent **(a)** depends on the unique combination of smile, gaze and proxemics it exhibits towards the subject during the first 12.5 seconds of the interaction and **(b)** is further moderated by the subject's own personality.
- **H2**: The amount of friendliness that subjects attribute to a greeting agent **(a)** depends on the unique combination of smile, gaze and proxemics it exhibits towards the subject during the first 12.5 seconds of the interaction and **(b)** is further moderated by the subject's own personality.

4.2.2 Apparatus and Stimuli

The context was a virtual main entrance of a museum. The scene always started with the subject's avatar (AVATAR) outside, in front of automatic sliding doors, and the greeting agent (AGENT) standing inside, close to a reception desk watching a computer screen. Figure 4.2 shows this setting in first (left) and third (right) person perspective when the approach has already started.

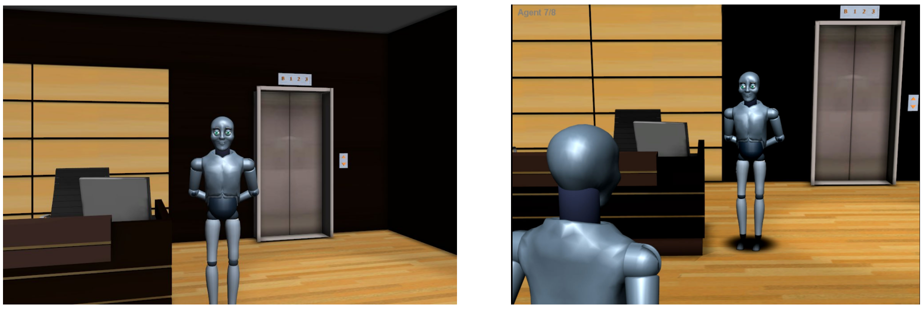

(a) First person view (b) Third person view

Figure 4.2: The setting of our study with the user's avatar entering the virtual museum entrance in first (left) and third (right) person mode and the greeting agent waiting inside.

To conduct the study in a fully controlled fashion and have subjects focusing exclusively on the AGENT, their level of interaction was limited to deciding when to start the approach by pressing a specific button. This triggered a locomotion behavior of their AVATAR towards the AGENT that automatically ended when the AVATAR reached the encounter space. We limited the control of the AVATAR to this simple choice to ensure that all approaches were performed in the same way across all conditions and subjects. To control for possible bias of the agent's visual appearance on the impressions formed, the agents were always graphically identical and not wearing any clothes. We used a male gendered model having human resemblance. Body movements were generated with procedural animation techniques and included a default eye blinking behavior and a slight body oscillation movement. All AGENTS were *always* holding the arms at the back with hands unclenched (as shown in Fig. 4.2). To give the idea of interaction with different entities we assigned them the name "Agent" followed by a progressing number shown at the beginning of each approach and in the top-left corner of the screen.

Our independent variables were **smile** (no vs. yes), **gaze** (low % vs. high %) and **proxemics** (no step vs. step). We conducted an informal manipulation check (N =

10, 2 females and 8 males, every subject tested both 1P and 3P perspectives) where we deployed a simplified version of the 3D environment and the agent exhibiting each behavior separately on both levels to verify that differences between the levels were correctly perceived by subjects within the duration of the approach (see Section A.1 in addendum for details).

The exact timing and location for triggering each behavior was based on Kendon's greeting model [Kendon, 1990] and Hall's proxemics theory [Hall, 1966]. Figure 4.3 shows a schematic top view of the scene with the AVATAR and the AGENT in their initial positions. The grayed dotted line shows the path followed by the AVATAR, black arcs are points where specific behaviors were exhibited. The description on top of them includes: a short reference name (in square brackets), the corresponding stage in Kendon's model (except for the custom point T2), the distance (in meters) from the AGENT and the name of the corresponding space in Hall's model (when overlapping). The arc without description was added to manipulate gaze (as described later) and the gray circular sections represents the AGENT's social and personal space according to Hall's proxemics theory. The duration of 12.5 seconds for each approach came naturally from the two models chosen: It was the time needed by the AVATAR to walk from its initial position (slightly off T1) to the *encounter* point (T4), that coincided with Hall's *personal space* boundary (humans usually do not allow others to cross this space, in particular during a first encounter). The duration was also determined by the AVATAR's speed, that was fine-tuned in the manipulation check to make sure that subjects were able to observe all the nonverbal cues exhibited by the AGENT, while keeping a walking speed for the AVATAR as much natural as possible.

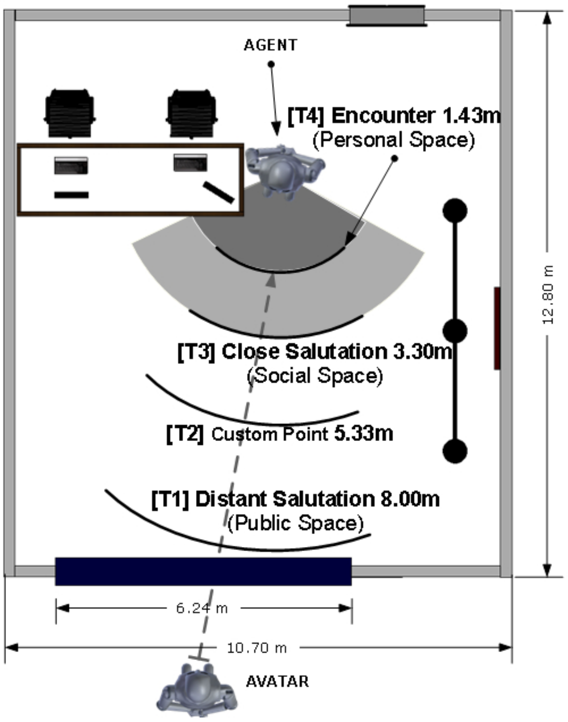

Figure 4.3: A schematic top view showing points as suggested by Kendon's model where specific behaviors were exhibited by the agent during the avatar's approach. Hall's proxemics zone names are shown in parenthesis.

We created a *baseline behavior* for the AGENT that was exhibited across all conditions of the study when the AVATAR approached it. This consisted of watching the computer screen at the beginning with both head and eyes towards it, gazing at the AVATAR for 2 seconds when it was at T1 (8m), looking back at the screen moving only the eyes and, finally, gazing at the AVATAR at T3 (3.30m). The AVATAR always stopped at T4 (1.43m). In the smiling condition the AGENT started smiling at T1. The *"high %"* gaze was obtained with a 2 seconds eye glance at T2. It follows that the difference between "low %" and "high %" gaze conditions was simply related to their duration, in the former the AGENT looked at the subject's AVATAR for a shorter time compared to the latter (in the manipulation check we validated whether subjects were able to distinguish between the two). The *"step"* condition was simply a step towards the AVATAR when it was at T3[1] keeping the arms at the back. Since we had eight different conditions, we adopted a latin square design to partially counter balance the treatment order and avoid first order carryover effects [Bradley, 1958].

4.2.3 Measures

A summary of our measures is provided in Table 4.1. ***Agent Extraversion*** was assessed using 4 items from the Saucier's Mini-Markers [Saucier, 1994] set of adjectives for measuring the Big 5. Two with positive (bold and extroverted) and two with negative (shy and withdrawn) valence. For the analysis the negative valence items scores were flipped and averaged with the positive ones to provide a final score. As exploratory variables, we included the ***Extra Impressions*** formed by subjects right after every approach and a measure of ***Agent Likeability***. The full set of questionnaire forms administered is listed in Section A.2.

4.2.4 Participants and Procedure

We had 32 participants for each trial recruited via public announcements in our university campus and the surrounding city. In the 1P trial we had 20 males and 12 females representing 11 nationalities[2]. In the 3P trial we had 19 males and 13 females representing 9 nationalities. In both trials, subjects were aged 21-60 with 63% in the 21-30 range. All subjects were well educated and most were

---

[1] Here the T2 in the published proceedings was a typo.

[2] As part of the demographic information, we asked participants to select the nation that most represented their cultural identity from a list of all countries in the world.

MEASURE	QUESTION	POINTS	LEFT ANCHOR	RIGHT ANCHOR
Agent Extraversion	I think the agent is [bold, extraverted, shy, withdrawn]	9	Extremely inaccurate	Extremely accurate
Agent Friendliness	How hostile/friendly has the agent been towards you?	9	Extremely hostile	Extremely friendly
Agent Likeability	Would you want to continue the interaction with this agent later?	5	No, definitely not	Yes, definitely
Subject Personality	Extraversion, agreeableness, neuroticism (using Saucier's items)			
Extra Impressions	Subjects asked to write adjectives that came to their minds			

Table 4.1: Summary of measures. Points refer to number of points on Likert scales.

at least familiar with computer science and psychology, detailed demographics are provided in the appendix section A.3. They were led to a dedicated room at our university facility, seated in front of a 19″ LCD monitor, instructed about the procedure and shown a tutorial for familiarization. After this introduction, the investigator monitored the session from an adjacent room. The session consisted of (1) observing each approach and then filling a form that included all measurements except the subject personality, (2) completing the personality inventory and (3) inserting demographic data in separate web forms. Finally, the investigator debriefed them. All the documents used are reported in Section A.4.

4.2.5 Quantitative Results

We conducted separate statistical analyses for the two trials, an informal comparison between the two is discussed in Section 4.2.7. For each trial, we conducted a mixed-design ANOVA for each measure (Agent Extraversion and Friendliness) with smile, gaze, and proxemics as within-subjects factors and subject extraversion, agreeableness and neuroticism as between-subjects factors.

The assessment of agent's personality and friendliness was based on the same scales made by [Saucier, 1994]. The usage of likert scales as interval measures and, therefore, the adoption of parametric statics to analyze them (i.e. ANOVA) is the subject of an ongoing debate [Jamieson, 2004]. It is suggested to treat them as ordinal measurements and use non-parametric methods for the data analysis.

However, as [Carifio and Perla, 2007] suggest, our rating scales constitute a fair approximation of interval level measurements since we have labeled our items suggesting the numeric association with a clear central marker, they have more than 5 points and yielded normally distributed data.

We used a full factorial model except that we omitted interactions among the between-subject factors. In order to use the three subject personality traits as between factors, for each measured trait we split our population in tertiles, thus resulting in 3 levels *"low, medium* and *high"* for each trait. For quantitative variables this has been shown to be a better practice [Gelman and Park, 2009] compared to the median split [MacCallum et al., 2002] (*"high"* and *"low"*). Main effects of interactions between factors are tested using Bonferroni adjustments for multiple comparisons. Table 4.2 provides a summary of our quantitative findings for both trials (see Section A.5 in appendix for the means and ANOVA results).

TRIAL	AGENT EXTRAVERSION	AGENT FRIENDLINESS
1P	Smile (.34)	**Smile (.000)**
	Gaze (.082)	**Gaze (.049)**
	Proxemics (.000)	Proxemics (.157)
	Gaze * S. Extraversion (.052)	Gaze * S. Agreeableness (.026)
	Smile * S. Agreeableness (.084)	Smile * Proxemics * S. Agreeable. (.03)
3P	Smile (.54)	**Smile (.000)**
	Gaze (.182)	**Gaze (.002)**
	Proxemics (.000)	Proxemics (.303)
	Smile * S. Extraversion (.025)	Smile * S. Extraversion (.002)
	Gaze * Proxemics * S. Extra. (.057)	Smile * S. Agreeableness (.064)
	Smile * Proxemics * S. Neuro. (.070)	Gaze * Proxemics * S. Agreeable. (.077)

Table 4.2: A summary of our results. The first column indicates the camera perspective of the trial, second and third refer to our two measures: agent extraversion and friendliness. For each measure relevant main effects and factor interactions, including significance level (p-values in parenthesis), are reported. All main effects positively affected extraversion and friendliness. The factor interactions had different influence depending on the subject personality. The abbreviation S. stands for *"subject"*.

First Person Perspective (1P)

Agent Extraversion. The analysis revealed a significant main effect of proxemics on agent's perceived level of extraversion, ($F(1,25) = 34.75, p < .001, \eta_p^2 = .58$); *stepping* agents were rated higher than *non-stepping* agents (**H1-a supported**). The main effect of gaze was near significant, ($F(1, 25) = 3.28, p = .082, \eta_p^2 = .12$). The main effect of smile was not significant ($F(1, 25) = .937, p = .34, \eta_p^2 = .036$), and there were no significant factor interaction effects. However, the factor interaction between gaze and subject extraversion was near significant, ($F(2, 25) = 3.35, p = .052, \eta_p^2 = .21$), as was the factor interaction between smile and subject agreeableness, ($F(2, 25) = 2.74, p = .084, \eta_p^2 = .18$), therefore **H1-b** was **rejected**.

Agent Friendliness. There was a significant main effect of smile on agent's perceived level of friendliness, ($F(1, 25) = 34.75, p < .001, \eta_p^2 = .58$); *smiling* agents were rated higher than *not smiling* ones (**H2-a supported**). There was a significant main effect of gaze, ($F(1, 25) = 4.27, p < .05, \eta_p^2 = .15$), and a significant factor interaction between gaze and subject agreeableness, ($F(2, 25) = 4.2, p < .05, \eta_p^2 = .25$). This would suggest that the effect of gaze depended on the subject personality. A main effects follow-up analysis revealed that gaze affected the ratings of agent friendliness for *low* agreeable subjects ($F(1, 25) = 10.44, p = .003$), but not *medium* ($F(1, 25) = .68, p = .41$) and *high* ones ($F(1, 25) = .63, p = .43$), thus **H2-b** was **partially supported**. The main effects of gaze were further analyzed by pairwise comparisons with Bonferroni adjustment for multiple comparisons. For subjects with *low* level of agreeableness, the ratings of agent friendliness in the *low gaze* condition ($M = 5.05, SE = .34$) were significantly lower than the *high gaze* condition ones ($M = 5.96, SE = .30$). There was also a significant factor interaction between smile, proxemics and subject agreeableness, ($F(2, 25) = 4.02, p < .05, \eta_p^2 = .24$). The follow-up analysis of proxemics main effects was not significant (all $F \leq 3.39, p \geq .077, \eta_p^2 \leq .12$). On the other hand, smile affected the ratings of agent friendliness at all levels of proxemics and for all the three subject personality levels (all $p \leq .015$), except for *low* agreeable subjects when the agents were *not approaching* them ($p = .36$). All other main effects and interactions were non-significant or irrelevant to our hypotheses (all $F \leq 2.68, p \geq .088, \eta_p^2 \leq .17$).

Agent Likeability. We ran the same mixed-design ANOVA for the ratings of agent likeability. There was a significant main effect of smile on agent's perceived likeability ($F(1, 25) = 20.03, p < .001, \eta_p^2 = .44$); subjects preferred to continue the interaction with *smiling* agents.

Third Person Perspective (3P)

Agent Extraversion. Results of the analysis revealed a significant main effect of proxemics on agent's perceived level of extraversion, ($F(1,25) = 67.20, p < .001, \eta_p^2 = .72$), and this was rated higher in the *step* condition (**H1-a supported**). The main effects of smile and gaze were not significant (Smile: $F(1,25) = .38, p = .542, \eta_p^2 = .015$; Gaze: $F(1,25) = 1.88, p = .182, \eta_p^2 = .07$). There was a significant factor interaction between smile and subject extraversion, ($F(2,25) = 4.27$, $p < .05, \eta_p^2 = .25$). This would suggest that the effect of smile depended on the subject personality. However, a main effects follow-up analysis revealed that smile affected the ratings of agent extraversion for *low* extraverted subjects ($F(1,25) = 5.90, p = .023$), but not *medium* ($F(1,25) = 3.05, p = .09$) and *high* ($F(1,25) = .14, p = .71$) ones (**H1-b** was **partially supported**). A main effects analysis indicated that for subjects with *low* level of extraversion the ratings of agent extraversion when *not smiling* ($M = 5.38, SE = .23$) were significantly different from the condition with *smiling* ($M = 5.84, SE = .23$). The factor interaction between gaze, proxemics and subject extraversion was near significant, ($F(2,25) = 3.29, p = .057, \eta_p^2 = .20$), as was the factor interaction between smile, proxemics and subject neuroticism, ($F(2,25) = 2.97, p = .070, \eta_p^2 = .19$). All other main effects and interactions were non-significant or irrelevant to our hypotheses (all $F \leq 2.06, p \geq .15, \eta_p^2 \leq .14$).

Agent Friendliness. There were significant main effects of smile and gaze on agent's perceived level of friendliness (Smile: $F(1,25) = 49.07, p < .001, \eta_p^2 = .66$; Gaze: $F(1,25) = 12.33, p < .005, \eta_p^2 = .33$); friendliness was rated higher either when the agent was *smiling* or when the *amount* of *gaze* was *high* (**H1-a supported**). The main effect of proxemics was not significant ($F(1,25) = 1.1, p = .303, \eta_p^2 = .04$). There was a significant factor interaction between smile and subject extraversion, ($F(2,25) = 8.0$, $p < .005, \eta_p^2 = .4$). This would suggest that the effect of smile depended on the subject personality. However, a main effects follow-up analysis revealed that smile affected the ratings of agent friendliness for *medium* ($F(1,25) = 49.94, p = .000$) and *high* extraverted ($F(1,25) = 9.74, p = .005$) subjects, but not *low* ones ($F(1,25) = 2.58, p = .12$), thus **H2-b** was **partially supported**. The main effects of smile were further analyzed: for subjects with *medium* level of extraversion the ratings of agent friendliness when *not smiling* ($M = 4.68, SE = .24$) were significantly lower than conditions with *smiling* agents ($M = 6.87, SE = .28$). For subjects with *high* level of extraversion the ratings of agent friendliness when *not smiling* ($M = 5.44, SE = .27$) were significantly different from the conditions with *smiling* ($M = 6.4, SE = .28$). The factor interaction between smile and subject agreeableness was near significant, $F(2,25) = 3.08$, $p = .064$, $\eta_p^2 = .19$, as was the factor interaction between gaze, proxemics and subject agreeableness ($F(2,25) = 2.85, p = .077, \eta_p^2 = .19$), and the interaction smile, gaze, proximity and subject agreeableness ($F(2,25) = 2.98, p = .069, \eta_p^2 = .19$). All other main effects and interactions were non-significant or irrelevant to our hypotheses (all $F \leq 2.53, p \geq .099, \eta_p^2 \leq .17$).

Agent Likeability. There were significant main effects of smile and gaze on agent's perceived likeability (Smile: $F(1,25) = 41.35, p < .001, \eta_p^2 = .62$; Gaze: $F(1,25) = 9.91$, $p < .005$, $\eta_p^2 = .28$); subjects preferred to continue the interaction with agents *smiling* and *gazing* at them *more*. The factor interaction between smile and subject extraversion was near significant, ($F(2,25) = 2.68, p = .088, \eta_p^2 = .17$), as was the factor interaction between proxemics and subject extraversion, ($F(2,25) = 2.73, p = .084, \eta_p^2 = .18$).

4.2.6 Qualitative Results

For the analysis of **"Extra Impressions"** we grouped synonymous adjectives into different categories. For each of these, we counted the number of different subjects that used adjectives belonging to that category. The full dataset already grouped in categories is provided in Section A.6 in addendum. Table 4.3 shows an excerpt of relevant categories of adjectives (with total counting for each trial) that we discuss in this qualitative analysis.

ADJECTIVES	TOTAL COUNT	
	1P	3P
Bored, annoyed, tired	24	15
Kind, polite, gentle	20	6
Authority, powerful, leader, achiever, ambitioned	17	10
Aggressive, stern, challenging and unfriendly	15	18
Careless, dismissive, uninterested	12	23
Professional, business-like, precise	12	10

Table 4.3: An excerpt of relevant adjectives discussed in this qualitative analysis of the subjects' extra impressions of the agents. The columns on the right show the total counts (i.e. the number of different subjects that used the adjectives belonging to that category) for each trial.

In both trials subjects' extra impressions revealed that the agent was judged as "bored, annoyed" mainly when *not smiling* and *not stepping* or exhibiting a *short* amount of *gaze* and judged as "careless, dismissive, uninterested" when *smiling* but *gazing* for a *short* amount of time and vice versa. Impressions of "aggressive, stern, challenging and unfriendly" were formed when the agents were *stepping*. In general, subjects judged the agents as "kind, polite, gentle" and used common human characteristics to define their extra impressions, thus perceiving the agents as believable even though all our behaviors were pre-scripted. Only a few subjects used adjectives such as "fake, deliberated, agent, scripted" (Tot. 1P = 2, 3P = 6) in the specific condition when he was *stepping*, *not smiling* and *gazing* briefly. Furthermore, adjectives such as "authority, powerful, leader, achiever, ambitioned" were used mainly when *stepping* and "professional, business-like, precise" when *not smiling* regardless of proxemics and gaze levels.

4.2.7 Discussion

For the first person perspective (1P), H1-a and H2-a were supported. We found that the amount of extraversion and friendliness that subjects attributed to our agents depended on unique combinations of smile, gaze and proxemics that they exhibited. In particular, agents *stepping* towards the subject's avatar were judged as more extraverted than agents *not stepping*, regardless of gaze amounts or whether they were *smiling* or *not*. Smile had a main effect on judgments of friendliness. These results seem quite intuitive but it is important to note that proxemics had absolutely no effects on judgments of friendliness even though qualitative impressions of "aggressive, stern, challenging and unfriendly" were formed when subjects judged *stepping* agents. Therefore, we had a sharp distinction between interpretations of proxemics and smile. When it came to judging extraversion proxemics had the highest weight, whereas smile dominated the impression formation of friendliness. This is an important result if we consider that smile and gaze can also be used to express personality traits (extraversion) as suggested by previous findings in human social psychology [Argyle, 1988, Campbell and Rushton, 1978, Mehu et al., 2008].

The relation between subject personality and behavior interpretation is harder to explain since H1-b was rejected and H2-b only partially supported. The effect of gaze on agent friendliness partially depended on subject agreeableness. *Low* agreeable subjects interpreted more gazing friendlier compared to less gazing. We didn't get significant results for *medium* and *high* agreeable subjects. According to the personality inventory we used, those who scored *low* in agreeableness are likely to be cold, unsympathetic, rude and harsh as opposed to the warm, kind and cooperative highly agreeable people. We think that this might reflect results of a previous study arguing that low sociable people tend to be more accurate in judging others in zero-acquaintance situations [Ambady et al., 1995]. The factor interaction between gaze and subject extraversion was near significant for the agent level of extraversion, and again only for *low* extraverted subjects (shy, quiet, withdrawn).

Gaze is also involved in a possible explanation for the factor interaction between smile, proxemics and subject agreeableness when judging the agent friendliness. In fact, smile had effect on all the subjects except the *low* agreeable group in the particular conditions when the agents were *not stepping*. This would suggest that this group gave more importance to gaze in that case. Although non-significant, a similar trend was observed also in the judgments of agent extraversion.

H1-a and H2-a were also supported when moving to third person perspective and with quite similar results. Again agents *stepping* towards the subject's avatar were judged as more extraverted than agents *not stepping* towards them, regardless of smile and the amount of gaze they gave. The effects of gaze on agent friendliness were clearer and didn't depend on subjects' personality. They interpreted agents gazing more at them as friendlier. Smile also led to higher ratings of friendliness, except for *low* extraverted subjects that formed impressions of extraversion rather than friendliness when judging a smiling agent. A possible explanation could be still related to the higher accuracy of judgments that low sociable people express, therefore interpreting smile as a cue of higher extraversion in that case. Another reason could be the great variability we had in the subjects level of extraversion (2.25 to 8.13) whereas the level of agreeableness was more compact (5.00 to 8.25). In general, the role of smile and proxemics was clearly separated also for this trial.

Our findings indicate that results in social psychology research on the assessment of personality traits and attitudes on the basis of nonverbal behavior [Burgoon et al., 1984, Riggio and Friedman, 1986, Argyle, 1988] do translate to the context of user-agent interaction. In particular, outcomes of using nonverbal immediacy [Richmond et al., 2008] are preserved in virtual encounters.

Proxemics behavior is of great importance for managing impression of extraversion. In order to manage impressions of friendliness, smile has the biggest weight, followed by gaze. Gaze, in particular, depends more on the user's personality, although this is not completely clear from our findings due to lack of significance. We think that variation of our subjects' cultural identity influenced this result. In particular we had greater variability in the 1P trial (53% Iceland, 9% Germany and 38% nine different nationalities), compared to 3P (75% Iceland and 25% eight different nationalities) and we think that this also influenced subjects' judgments making hard to get significant results in the personality interactions. However, from our statistical analysis we can claim that the neuroticism trait doesn't have any effect when interpreting the nonverbal cues we manipulated. We are also aware of the existence of gender differences in interpretation of nonverbal cues both in human-human [Argyle, 1988, Miller et al., 2007, Richmond et al., 2008] and human-agent interactions [Kulms et al., 2011].

Despite a stronger effect of gaze in 3P, results in both trials were similar, thus suggesting that camera perspective does not alter the way our set of nonverbal cues was interpreted. This result reflects our expectations, even though we couldn't formulate a precise hypothesis a priori due to the lack of previous work investigating this particular aspect. Similar research dealt more with immersive virtual environments [Mohler et al., 2008] explored with head mounted displays [Salamin et al., 2006, Mohler et al., 2010] but not with 3D virtual environments experienced in the same way as in our study or in many of the works mentioned in chapter 3. We think that, in addition to impacting the virtual agents community, this result has also implications in the study of human social psychology. It is interesting to see how users in the 3P trial were still able to form impressions of a virtual character (the agent) when it was exhibiting nonverbal cues towards another virtual character shown on the screen (their avatar) and not directly towards them as in 1P, thus putting themselves completely in the role of a virtual entity external to their body.

Furthermore, in both trials results of agent likeability mirrored those of friendliness, thus agents smiling and gazing more also resulted in more approachable and likable agents. This is not surprising considering that one of the advantages of immediacy cues is obtaining a more favorable impression [Richmond et al., 2008], but it also foresees that friendliness was considered more important than extraversion by subjects when they had to decide whether to continue the interaction or not.

Some limitations should be considered. When we looked at the relationship between subjects' personality and their interpretations we found interesting trends supporting that personality acted as moderator. However, these speculations are limited by the statistical significance of the results and the specific population obtained. The ideal body of subjects would have consisted of a balanced population with personality equally distributed in the three groups for each trait. Furthermore, we are aware that cultural identity has influence on behavior interpretation and, in particular, in the 1P we had a high variety in the population.

4.2.8 Conclusion

This study has provided evidence that people form impressions of a greeting agent's extraversion and friendliness in the very first moments of a virtual encounter. In our experiment, agents exclusively exhibited nonverbal immediacy cues (smile, gaze, proxemics) when users approached them with their own avatar in a reception of a virtual museum.

We also discovered that the specific cues adopted had well defined and separate meanings along the two dimensions measured, in particular proxemic behavior had the biggest impact on impressions of extraversion, whereas smile and gaze had significant impact on friendliness.

We demonstrated that camera perspective (first or third view) has no influence on user interpretations of an agent's nonverbal behavior in this context.

Finally, **"Agent Likeability"** results mirrored those of **"Agent Friendliness"**, smiling and gazing agents were considered more approachable and likable. According to our measurement of likability, these results foresee that subjects might consider friendliness (attitude) more important than extraversion (personality) when it comes to deciding whether to continue the interaction or not later. The study described in the following section was built on these findings and aimed at understanding the impact that these two first impression dimensions might have on users' relational behavior towards a virtual agent.

4.3 Nonverbal Behavior Impact Study

This study aimed at investigating whether users' first impressions of an agent has impact on their relational decisions in terms of how likely they are to interact with the agent again and how often they would then interact with it. We also wanted to study possible effects of concordance between the user′s own personality and the agent's personality and attitude when making these decisions, i.e. whether the user′s relational decisions depend on his/her own personality as well.

Subjects observed a series of animated views of first-person perspective approaches to life-sized agents presented as guides of a virtual museum. The guides exhibited two levels of personality (high vs. low extraversion) and attitude (high vs. low friendliness) towards the subjects during initial greeting approaches of 12.5 seconds. The manipulations, described later in Section 4.3.1, were exclusively based on nonverbal immediacy cues exhibited by the guides as suggested by our previous study. After meeting each guide, subjects were immediately asked to express, in general, how likely they would be to spend time with it again and their preferences about the number of guided tours they would be willing to take with the agent. Details about the measurements follow in Section 4.3.2, and in Section 4.3.3 we explain the experimental procedure in more detail.

4.3.1 Apparatus and Stimuli

The subjects interacted with our guides standing still in front of a 2.7mx2.0m (width x height) screen at a distance of 1.50m. Figure 4.4 shows a 3D representation of this setting. The screen projected life-size agents in a 3D virtual entrance of a museum. The scene always started with the subject outside, in front of automatic sliding doors and a single guide agent waiting inside, standing close to a reception desk, holding its arms behind its back and watching a computer screen, similar to our previous study.

From the subject's point of view, an experimental session consisted of observing (in first person perspective) their approach to a series of virtual guides to symbolically give each of them a registration card, as shown in Figure 4.4.

A touch screen tablet computer was used to give subjects the ability to start each interaction, by pressing a button on the screen, and to administer self-report questionnaires at the end of each interaction. When a subject touched the start button on the tablet, it triggered an animated locomotion towards the guide agent. It automatically stopped when the subject reached the encounter space of the agent. The self-report questions were displayed after approaching each guide and subjects tapped on the tablet screen to reply.

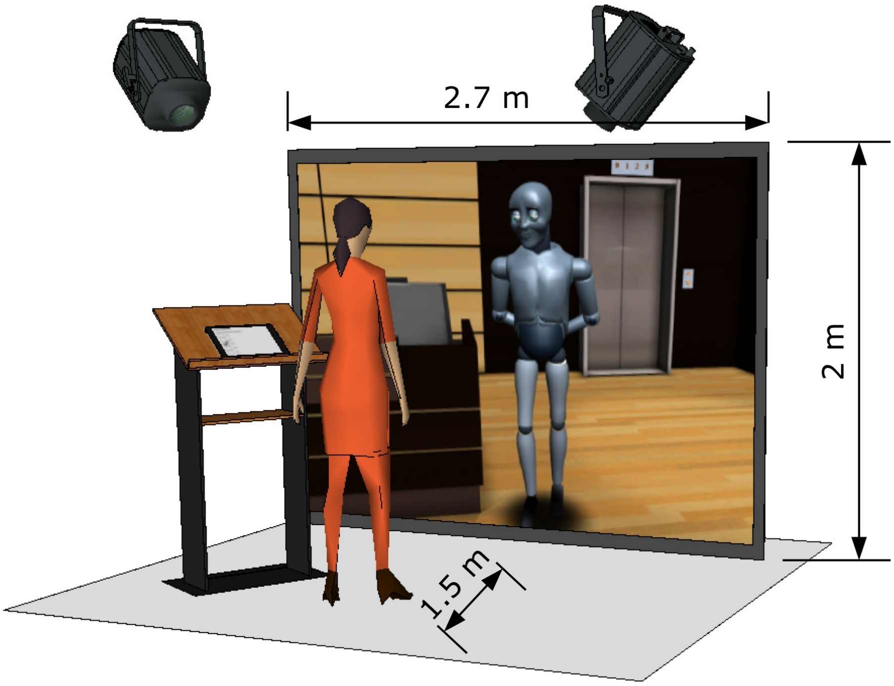

Figure 4.4: A 3D reconstruction of the setting adopted for our study showing the user position standing still in front of an animated view of the full-size guides. The tablet computer on the kiosk at the left of the user was used to answer the questions after meeting each guide in the virtual museum entrance shown on the screen in front of the user.

A system of lights in the ceiling of the experiment room were controlled via software to lead the subject through the various steps of the study. The purpose was to show always "where to stand". Therefore, a light would turn on in front of the screen to make a marker visible on the floor. This marker indicated the point where the subject had to stand when observing the approach towards the guide on the screen. Another light would turn on (switching the other off) to show that questions were available on the tablet right after approaching each guide. Pictures of the experimental room showing this setting are provided in appendix B.1.

The independent variables were **Guide Extraversion** (low vs. high) and **Guide Friendliness** (low vs. high). The different levels of extraversion and friendliness were obtained by manipulating agents' nonverbal immediacy cues of smile, gaze and proxemics according to results of our previous study. Since we discovered that agents that stepped towards approaching users were judged as more extraverted, our guides did the same when **guide extraversion** was at *high* level, and did not step towards the subjects when at *low*. On the other hand, we earlier found that smiling and gazing agents were judged as more friendly, therefore we had smiling and more gazing guides for the *high* level of **guide friendliness**. Table 4.4 shows all the mappings between the levels of our IVs and the resulting nonverbal behavior exhibited by the guides. To control for ordering and carryover effects, we showed the resulting four different conditions with a full counterbalanced treatment order [Bordens and Abbott, 2002] in a within-subjects design.

GUIDE level of		NONVERBAL CUES
EXTRAVERSION	FRIENDLINESS	
Low	Low	No smile, low % gaze, no step
Low	High	Smile, high % gaze, no step
High	Low	No smile, low % gaze, step
High	High	Smile, high % gaze, step

Table 4.4: Mapping of our IVs to the nonverbal cues exhibited by the guides according to our previous findings in the Nonverbal Behavior Interpretation Study.

4.3.2 Measures

Table 4.5 provides a summary of our measures. The first two measures assessed relational decisions: ***Likelihood of Encounters*** assessed the overall likelihood of future guided visits, while ***Number of Visits*** assessed decisions about the frequency of those visits. After meeting the four guides, we asked subjects to choose their preferred guide (***Guide Preference***) among the four that they met and we assessed their personality (***Subject Personality***) using the Saucier's Mini-Markers [Saucier, 1994] set of adjectives for extraversion and agreeableness traits. The full set of questionnaire form administered is listed in Section B.2 in appendix.

MEASURE	QUESTION	RANGE
Likelihood of Encounters	Would you like to do business with this guide?	1 (No, definitely not) 5 (Yes, definitely)
Number of Visits	How many guided visits would you like to take with this guide?	0-10
Guide Preference	Of the four guides you have just met, which one would you prefer as your guide?	Guide [1-4] or "I don't know"
Subject Personality	Extraversion and Agreeableness traits	

Table 4.5: Summary of measures. Likelihood of Encounters is on a 5-point scale.

4.3.3 Participants and Procedure

We had 24 participants recruited via public announcements in our university campus and the surrounding city. There were 15 males and 9 females representing 8 nationalities[3]. Subjects were aged 21-60 with 54% in the 21-30 range. All subjects were well educated and most were at least familiar with computer science and psychology. Detailed demographic data are provided in Section B.3. They were led to a dedicated room at our university facility, instructed about the procedure, shown a tutorial for familiarization and asked to sign the consent declaration. Subjects were led to believe that the selected number of visits represented an actual time commitment. The following is an excerpt of the consent form that we used:

> *"By signing this document, I agree to come back in these facilities and be guided in one or more virtual tours of the museum according to the particular guide that will be assigned to me and the preference (number of visits) that I expressed for him. The guided tours (if any) will be scheduled over a period of two consecutive months at my earliest convenience. The start date of this period and all appointments will be scheduled in concordance with the investigator. Every visit will require approximately 15 minutes and different area of the museum will be shown if more than one visit is scheduled."*

---

[3] As part of the demographic information, we asked participants to optionally select the nation that most represented their cultural identity from a list of all countries in the world as in the behavior interpretation study, only one subject didn't choose any.

We explained to our subjects that the assignment of the guide would be random, but we would respect their wishes regarding each one (i.e. how often they would like to see them again). This was introduced to prevent subjects expressing realistic preferences for only the most preferred guide and giving zero or a low number of visits to the others. After signing the consent declaration, we clarified again that a tour would require them to come back to our facilities and spend time with the guide delivering it, then the investigator left the room to monitor the session from an adjacent observation room. The session consisted of (1) meeting each guide and then replying to questions on the tablet including the likelihood of future encounters and the number of visits, (2) choosing the preferred guide, (3) completing the personality inventory and (3) providing demographic data. Finally, the investigator debriefed them. For more details about the documents used for this study see Section B.4 in appendix.

The first five subjects of our study were quickly interviewed after the debriefing. Our goal was to check whether (a) they believed in the real time commitment that they were taking by giving different options for the number of visits, (b) the choice was not influenced by their time availability and (c) the virtual museum realm and the possibility of having guide tours was of interest or not for them or, in general, biasing their decisions. Only one subject reported having limited time availability in the near future, but all the others stated that their choices were only influenced by the guides presented. They all claimed that they seriously considered the possibility of coming back for subsequent visits and that the museum scenario had relatively little effect on their decisions compared to the guides' behavior.

4.3.4 Hypotheses

From the results of our previous study and findings in social psychology on the outcome of nonverbal immediacy [Richmond et al., 2008, Burgoon et al., 1984], we predicted that **(a)** friendliness, as an interpersonal attitude and a "short-term" feature, would have a main effect for all our measurements, whereas **(b)** extraversion, as a personality trait of the guides and a "long-term" feature, would interact with the subject's own personality. Although prior results on the effect of user-agent personality matching are inconsistent (see Section 3.4), we predicted that similarity in the agent's personality (**Guide Extraversion**) and the user's personality would have positive effects on outcomes. These were our hypotheses:

- **H1 (Likelihood of Encounters)**: **(a)** Guide Friendliness will have a main effect on the likelihood of future encounters. Subjects will be more likely to encounter *high* friendliness guides again compared to the *low* friendliness ones; **(b)** Subjects own personality, either the extraversion or the agreeableness trait, will positively interact with guide extraversion (e.g.: *High* extraverted subjects will be more likely to encounter *high* extraverted guides again);
- **H2 (Number of Visits)**: **(a)** Guide Friendliness will have a main effect on the number of future visits. Subjects will choose a higher number of visits with *high* friendliness guides compared to the *low* friendliness ones; **(b)** Subjects own personality, either the extraversion or the agreeableness trait, will positively interact with guide extraversion (e.g.: *High* extraverted subjects will choose a higher number of visits with *high* extraverted guides);
- **H3 (Guide Preference)**: The preferred guide will have *high* level of friendliness.

4.3.5 Results

We conducted separate mixed-design ANOVAs for the two dependent variables ***Likelihood of Encounters*** and ***Number of Visits***, with guide extraversion and guide friendliness as within-subjects factors and subject extraversion and agreeableness as between-subjects factors. We used a factorial model, but we omitted interactions among the between-subject factors. In order to use the two subject personality traits as between-subjects factors, for each measured trait we split our population at the median into *"low"* and *"high"* groups. Details about the means and ANOVA results can be found in Section B.5 in appendix.

Likelihood of Encounters

A mixed-design ANOVA was conducted for this measure. The analysis revealed a significant main effect of guide friendliness, $F(1,21) = 21.91, p = .000$, $\eta_p^2 = .51$; subjects were more likely to do business with *high* friendliness guides later, compared to the *low* friendliness ones (**H1-a supported**). The main effect of guide extraversion was not significant ($F(1,21) = .39, p = .54, \eta_p^2 = .018$). No significant interaction effects were observed between the guide extraversion and subject personality traits, therefore **H1-b** was **rejected**.

Number of Visits

This variable was measured on an ordinal scale ranging from 0 to 10, and did not follow a normal distribution. In order to use standard analysis techniques for repeated measures (i.e. RM-ANOVA), we applied an aligned rank transform (ART) for non-parametric factorial data analysis as suggested by [Wobbrock et al., 2011]. The analysis of the transformed values revealed a significant main effect of guide friendliness on the number of visits, $F(1, 20) = 14.22, p \leq .001, \eta_p^2 = .416$; subjects preferred to take a higher number of visits with *high* friendliness guides compared to the *low* friendliness ones (**H2-a supported**). The main effect of guide extraversion was not significant ($F(1, 20) = .062, p = .80, \eta_p^2 = .003$). No significant interaction effects between Guide Extraversion and the subject personality traits were observed, therefore **H2-b** was **rejected**.

Guide Preference

All the subjects, except one, had a preference for a specific guide type among the four presented as shown in Table 4.6. Subjects showed a highly significant preference for the guides with a *high* level of friendliness, $\chi^2(2, N = 23) = 12.56, p < .001$, therefore **H3** is **supported**. For the chi-square test of goodness-of-fit that was performed we dropped the single subject without preference and grouped the other preferences in two categories depending on the guide's level of friendliness.

TYPE	GUIDE level of		PREFS
	EXTRAVERSION	FRIENDLINESS	
1	Low	Low	3
2	Low	High	10
3	High	Low	0
4	High	High	10
5	"I don't know"		1

Table 4.6: The number of preferences (PREFS) received by each of the four types of guides presented in our study.

As an exploratory analysis, we shifted our focus to the preferred level of guide extraversion ignoring guide friendliness. To simplify this analysis, again we dropped the single subject without preference, thus resulting in a new dependent variable named **Preferred Guide Extraversion** with two levels: *low* and *high*. A logistic regression analysis was conducted to predict the preferred guide's level of extraversion for 23 subjects using the subject extraversion and agreeableness traits as predictors. For each of the personality traits measured, we used standardized z-scores of the original continuous values that were on the scale [1-9]. A test of the full model against a constant only model was statistically significant ($\chi^2 = 9.40, p < .05, df = 3$). Nagelkerke's R^2 of .450 indicated a moderate (not strong) relationship between guide extraversion and our predictors (subject's extraversion and agreeableness). Prediction success overall was 69.6% (76.9% for *low* extraversion and 60% for *high*). The Wald test demonstrated that there was only a near-significant interaction effect between Subject Extraversion and Subject Agreeableness on the preferred guide extraversion level (p = .071). These results were inconclusive: we observed a (near-significant) trend showing that subjects having a mismatch in their personality scores (high extraversion and low agreeableness, or vice versa) preferred guides with *high* level of extraversion, whereas subjects having a match tended to prefer guides with *low* level of extraversion. There were no significant main effects of the predictors.

Frequency Analysis for Number of Visits

In support of earlier results and to examine this measure from a different angle, we treated the number of visits preferences received by each guide as ordered categories and computed the frequencies of choices for every single category. Figure 4.5 shows histograms grouped by guide type (i.e. low vs. high levels of extraversion and friendliness). Note that no subjects selected 9 or 10 visits for any guide, and these options are not shown in the figure. For high friendliness guides — those most often preferred — subjects more frequently selected a high number of visits (5–8), compared to low friendliness guides where a low number of visits (0–1) were more commonly selected. The modes for the high friendliness guides were 3 visits (at *low* extraversion) and 2 visits (at *high* extraversion): A relatively high number of subjects would have spent from 30 (2x15) up to 45 (3x15) minutes with them later, considering that each visit would have required 15 minutes.

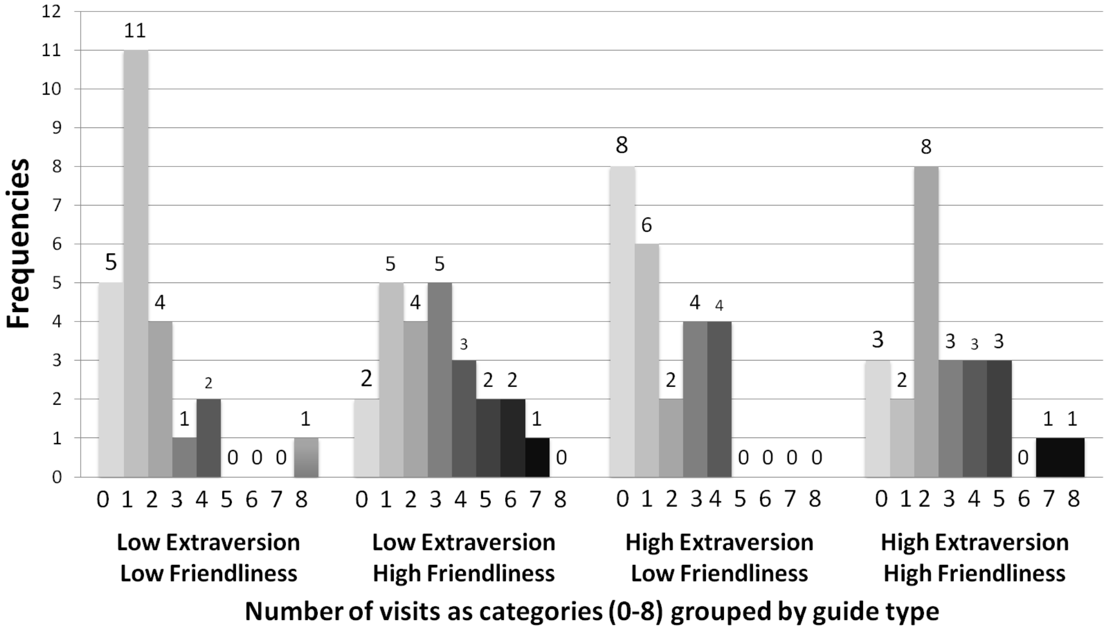

Figure 4.5: Histograms grouped by guide type showing: on the x-axis the ranges of options that subjects had when selecting the number of visits and on the y-axis the frequencies of selections for each option.

4.3.6 Discussion

Our major finding is that the impressions of our virtual guides, exhibiting the proper nonverbal immediacy cues at specific points during only 12.5 seconds of a first greeting encounter, had significant effects on users' relational decisions in terms of choosing how likely and for how often they would spend time with the guides later. The friendliness of the guide had a main effect for all our measurements, thus supporting our hypotheses (H1-a, H2-a and H3). In particular, *high* friendliness received a higher number of visit preferences than *low* friendliness guides, subjects wished more to do business with them later, and they were the most preferred ones regardless of the level of personality that was exhibited.

We think that this result has an important impact in the relational agents community and, in particular, in the museum agents application domain. Our findings suggest that Sunnafrank's theory [Sunnafrank, 1986] (described in Section 2.2) can be impactful for relational agents as well. In addition they do translate from human social psychology to human-agents encounters and demonstrate that impressions and relational assessments are made quickly, often in the beginning moments of initial conversations or, in our case, even prior to a conversation [Sunnafrank and Ramirez, 2004]. Furthermore, it is remarkable that all subjects, except for one, upon agreeing to come back in the consent, were able to choose their preferred guide among the four presented even though the interaction time was short, they had never seen our virtual guides before, and they also had the option to not choose any of them.

Contrary to what we expected, we did not observe any interaction effect between the user's own personality and the guide's level of extraversion, when measuring the number of visits and the likelihood of encounters (H1-b and H2-b rejected). A possible explanation is that considering the short amount of time that subjects had to form impressions of the guides, they based their decisions on a feature, interpersonal attitude, that happens in the "short-term" and is specifically addressed for them in that particular context, thus fulfilling their expectations about the guides being approachable, friendly and polite. As Argyle suggests, sometimes people are trying to establish a relationship with another person by being friendly for making friends or showing politeness during first encounters [Argyle, 1988]. So, this might have been the role that our guides assumed for the subjects, and it is not surprising that people would have these expectations, considering that in everyday interactions we always perform and adhere to conventions that have meaning to the audience, according to Goffman's dramaturgical study of social interaction in terms of theatrical performance [Goffman, 1959].

Furthermore, there is very little difference between the affiliation dimension of interpersonal attitudes and the agreeableness personality trait in terms of expressiveness through nonverbal behavior only [Richmond et al., 2008, Argyle, 1988]. There is evidence that agreeableness might be related to formation of alliances and cooperative relationships [Mehu et al., 2008]. Thus, we could have expected to have such interaction effects by focusing on manipulating the guide agreeableness level, instead of extraversion.

The use of our set of immediacy cues and the specific mapping adopted to express the guides' personality and attitude suggests another possible explanation. On one hand, the use of those cues altogether allows the agent to obtain a more favorable impression [Richmond et al., 2008], but on the other hand we should be aware that smiles are very potent interpersonal cues [DePaulo, 1992] that might have predominated over the others, leading to greater effects of friendliness [Mehu et al., 2008]. However, from these results we have learned that impressions of friendliness, in a virtual greeting encounter, can be easily obtained by properly manipulating the smile and gaze behavior of an agent according to Kendon's communicative model [Kendon, 1990]. Thus, we believe that some deployments of virtual agents can benefit from these results. For example, in museum exhibits such as [Bickmore et al., 2013, Traum et al., 2012], where there is limited space and thousands of people passing by, it might be advantageous not having to cope with proxemic cues that might be difficult to interpret by the multitude of users present.

Our subjects clearly preferred (20 out of 24) a friendly guide when asked to choose one among the four guides they met, therefore our interest moved to which level of extraversion subjects preferred the most. We will not speculate on the results of the exploratory analysis shown in Section 4.3.5, since we are aware that the significance levels obtained are not highly accurate for the type of analysis conducted (logistic regression), possibly due to the limited sample size (24 subjects). Future work is needed to investigate the role of the user's personality in terms of relational decisions and personality of the virtual agents presented.

4.3.7 Conclusion

In this study we found that the nonverbal behavior exhibited by our guides, in 12.5 second encounters, had significant effects on subjects' relational decisions in terms of how likely and how often they would like to spend time with the guides on virtual tours. In particular, guide friendliness, expressed with smiling and gazing more at the subjects, had a main effect for all our measures. We also found that personality concordance with guides had no significant effect on relational decisions.

Some limitations should be considered. First, our main outcome (number of visits) is a self-report measure and, even though our subjects consented to return to our lab according to the preferences they expressed, it was still a hypothetical decision. A behavioral measure, for example assessing whether subjects actually return for the tours they agreed to, would be preferable. This might be done within a longitudinal study design that would also allow us to examine the stability and long-term impact of attitudes and impressions formed in these brief user-agent encounters.

Secondly, the subjects' interactions with our guides were limited to observing the animated approaches in a 3D environment while they were physically standing still, whereas in many settings such as the museum agents deployed in public spaces described in [Bickmore et al., 2013, Gockley et al., 2005, Kopp et al., 2005, Traum et al., 2012, Lane et al., 2011], a more natural interaction that copes with visitors' physical proximity and their movements when they are walking around is required. In such scenarios, a multitude of visitors might approach an agent installation and quickly decide if it is worth their time by going ahead and engage in an interaction. In the following section we describe the last study of our series focusing on this particular scenario.

4.4 Managing First Impressions in a Public Space Study

Our intent was to test the effectiveness of managing first impressions of a virtual relational agent in a real application setting and give more freedom to the user when interacting with our agent. For this study we applied findings of behavior interpretation and impact study to Tinker [Bickmore et al., 2013], a museum agent guide installed at Boston Museum of Science[4]. The idea was to detect in real-time visitors approaching the installation and have the agent exhibit different affiliation (friendliness) levels as they were getting closer.

In particular, we conducted a fully randomized experiment with a three-group, between-subjects design, in which we had a control group with Tinker not exhibiting any nonverbal reaction towards an approaching visitor, then we compared a FRIENDLY TINKER version to an HOSTILE one by manipulating its greeting nonverbal behavior (i.e. by using smiling and gazing cues according to earlier findings) as described in the following section.

The goal was to evaluate whether visitors forming friendly impressions of the agent during the approach (a) were more attracted by and encouraged to engage in an interaction, and (b) were spending more time with the agent thereafter.

4.4.1 Apparatus and Stimuli

Tinker: the Relational Agent Museum Guide

Tinker is virtual museum guide agent developed by the Relational Agents Group at Northeastern University and currently installed in the Computer Place exhibit at the Boston Museum of Science, MA, USA.

The agent is projected onto a 46″ LCD tall display and it appears in the form of a human-sized anthropomorphic female robot as depicted in Figure 4.6. Tinker uses nonverbal conversational behavior, empathy, social dialogue, reciprocal self-disclosure and other relational behavior to establish social bonds with museum visitors. It uses a biometric identification system (based on hand image features) so that it can re-identify visitors it has already talked to, and maintains persistent discourse and relational models, so that prior conversations can be seamlessly continued. It also uses a multiple choice touch screen input for user utterances,

---

[4] See the web page at: `http://www.mos.org`

and to enable visitors to input their given name and to quickly jump to different high-level topics using iconic representations. The Figure 4.7 in the next section shows a 3D reconstruction of this setup.

In addition to a range of nonverbal conversational behavior –including hand gestures, head nods, gaze shifts, eyebrow raises, posture shifts, facial displays of emotion, and visemes (for lip synchronization with speech)– animations were developed so that Tinker could beckon approaching visitors, demonstrate how to use the hand reader, and sleep during idle periods.

Tinker is able to provide visitors with information on and directions to a range of exhibits in the museum, as well as discuss the theory and implementation underlying her own creation. The exhibit has been operational since September, 2007, and has already been seen by thousands of museum visitors [Bickmore et al., 2013].

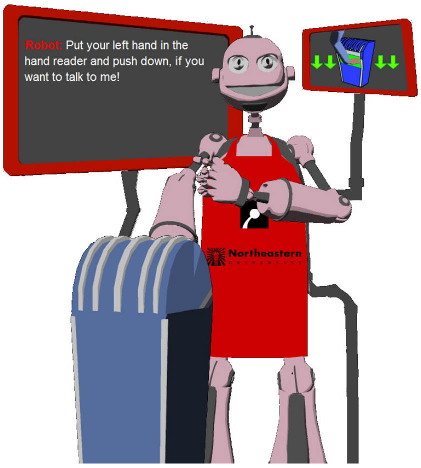

Robot: Put your left hand in the hand reader and push down, if you want to talk to me!

Figure 4.6: Tinker character appears as a female robotic agent. She wears the apron as a Museum of Science's Computer Place staff member. The large scrolling text screen placed behind Tinker shows the content of the last several conversational turns, the smaller sign on the right displays system status information and a demonstration animation sequence showing approaching visitors how to use the hand reader. The virtual hand recognition reader placed in front of Tinker is used by the agent to demonstrate to the visitors how to put their hand in the reader to start the conversation.

Experimental Conditions

Tinker originally had a generic motion sensor to determine if visitors were in its surroundings so that it could beckon them over to start a conversation. In order to implement our greeting model we needed better accuracy in detecting visitors proximity. For this reason we installed a Microsoft's Kinect sensor to detect their distance from the agent in real-time while approaching.

Figure 4.7 shows a 3D reconstruction of the area of the museum where Tinker is exhibited. Due to space constraints and kinect sensing capabilities, we needed to scale down by a **0.5 factor** the original distances as suggested by Kendon's greeting model and adopted in our previous studies. This allowed having the *"Distant Salutation"* within a distance of 4 meters from Tinker (instead of 8 meters), which was the best compromise between the least distance where this reaction could have happened and the constraints due to the spatial arrangement of Tinker in the room. As we can see in picture 4.7, Tinker lays in a corner facing the main entrance of the Computer Place area (down left) and a corridor that leads to different exhibits (down right), both at distance of about 4 meters. We kept the distance T4 as the original (instead of scaling it down to 0.78 meters). This was not possible due to kinect sensing limitations (it stopped sensing users that were less than 1 meter away). Furthermore, the position in which visitors had access to the input devices (hand reader and touch screen) was at about 1.40 meters as well. Pictures of the Computer Place space were Tinker is exhibited are provided in appendix C.1.

The black arcs in Figure 4.7 are points, according to our model, where Tinker's nonverbal reaction were triggered while the visitor was approaching. The description on top of the arcs includes: a short reference name (in square brackets), the corresponding stage in Kendon's model (except for the custom point T2), the distance (in meters) from Tinker and, in parenthesis, the original distance adopted in the previous studies. The arc without description was added to manipulate Tinker's gaze behavior for the FRIENDLY TINKER condition as described later.

The nonverbal immediacy cues exhibited by the agent were smiling and gazing behavior. According to our previous findings, smiling and gazing more at the user yielded for impressions of high affiliation (i.e. a friendly attitude), whereas not smiling and a lower amount of eye contact were judged as low affiliation cues (i.e. an hostile attitude). Table 4.7 shows a summary of the three study conditions and the corresponding nonverbal greeting behavior exhibited by Tinker.

In all conditions, Tinker beckoned the visitor when it reached the *"Encounter"* point T4 (1.4m) by demonstrating how to use the hand reader and inviting to begin the interaction. We had a CONTROL GROUP with Tinker not reacting at all during the approach. In this condition, Tinker was asleep having the arms crossed and the head a bit tilted on the right side. In both FRIENDLY and

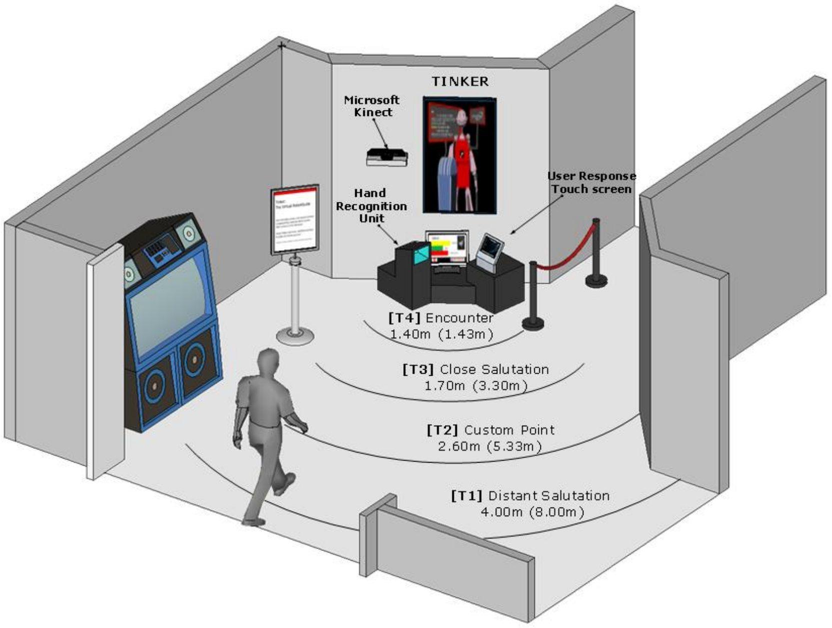

Figure 4.7: A 3D reconstruction of the Computer Place exhibit's entrance at the Boston Museum of Science showing Tinker's setup and points as suggested by Kendon's model where specific behaviors were exhibited by the agent during the visitor's approach.

HOSTILE TINKER conditions the agent waked up at T1 with a gaze at the visitor of 0.5 seconds. Then it looked away (on the right side) and looked back at the visitor when it was at T3 (1.7m). In the FRIENDLY TINKER condition we added a smiling facial expression at T1 (4m) and the *"high %"* of gaze behavior was obtained with a 0.5 seconds eye glance at T2 (2.6m). The gazes at the visitor had a shorter duration (0.5s) compared to the original one (2s) used in the Behavior Interpretation Study (see Section 4.2.2). This allowed Tinker to properly exhibit the gazes at the user given that we scaled down the distances to trigger the behavior.

Tinker nonverbal behaviors were all generated with pre-made animations. We conducted an informal manipulation check in our lab facilities (N = 10, 3 females and 7 males) with a deployment of Tinker similar to the exhibit at the museum but only with the Kinect sensor in operation and a similar flat LCD screen to display the agent. We asked subjects to approach Tinker while we were testing each behavior (i.e. smile and gaze) separately on both levels (neutral vs. smiling

Experimental Conditions	
GROUP	**MAPPED NONVERBAL CUES**
CONTROL	No reaction, Tinker wakes up from sleep
HOSTILE TINKER	Neutral face, low % of gaze at visitor
FRIENDLY TINKER	Smiling face, high % of gaze at visitor

Table 4.7: The three groups in our between-subjects study and the corresponding nonverbal cues exhibited by Tinker.

facial expression and low % vs. high % of gaze) to verify that differences between the levels were correctly perceived by subjects (see Section C.2 in addendum for details).

4.4.2 Measures

A summary of our measures is provided in Table 4.8. As Kendon suggests, how far one goes out of one's way to meet another appears to have a communicational significance [Kendon, 1990] and, most importantly, manifests the desire to engage in an interaction.

The ***Visitor Actions*** outcome measures how far visitors go when approaching Tinker. It is divided up in four possible outcomes that can be sequentially performed by a visitor and eventually lead to a conversation session with the agent. Visitors quitting (i.e. walking away) *at any point* during the approach perform no actions (i.e. *zero*). Those completing the approach, thus arriving at the *"Encounter"* point T4, are performing one action. At that point Tinker invites the visitor to begin the session by placing the hand in the reader, thus the second action is performed when the visitor puts the hand in the reader. Finally, after listening to Tinker's instructions on how to keep the hand in the reader and choose dialogue options from the touch screen input, the visitor can begin the conversation session with the agent. The very moment when the visitor actually begins the session determines the accomplishment of the last possible action that we measured.

The ***Session Duration*** is a conditional outcome depending on the number of actions taken before. It measures the duration of conversation sessions in seconds, however those are measured only for visitors that completed all of the actions prior to begin a conversation with Tinker.

MEASURE	DESCRIPTION	RANGE
Visitor Actions	Actions completed by the visitor before engaging in a conversation session with Tinker: 0. Visitor walks away; 1. Approach completed at T4; 2. Hand in the reader; 3. Conversation session can start.	[0 - 3]
Session Duration	Duration (in seconds) of a visitor conversation session with Tinker	[0 - ∞]

Table 4.8: Summary of measures. The actions taken by an approaching visitor leading to a conversation with Tinker are divided up in 4 possible outcomes. Visitors quitting while approaching (i.e. waling away) are performing zero actions. Once the approach is completed at T4 (one action), the visitor can place the hand in the reader (two actions) and after listening Tinker's instructions the conversation begins (three actions). The second outcome is conditional to the number of sub actions accomplished earlier and measures the duration of those conversations.

4.4.3 Participants and Procedure

It was impractical to obtain consent from visitors who approached but then continued past the exhibit, especially given the 5,000 visitors a month who walk by. Furthermore, asking for consent prior to beginning the approach would have primed the subjects invalidating the effects desired given that we were dealing with Tinker's first impressions. For these reasons, we anonymously collected only data about ***Number of Actions*** during visitors approaches and their ***Session Duration*** as purely anonymous behavioral outcomes.

After two weeks of pilot testing we collected data from the 11th of December 2012 to the 15th of October 2013 (309 days). Given the nature of the study, we couldn't ask for specific demographic information, but to give a general idea of the visitors that checked in, we report here a sample of the general visitors' demographics for the museum. This data was obtained from the surveys administered by the museum staff to 1.233 visitors from July 2012 to June 2013.

Visitors' age was mainly within the 35–54 range (57%). The 77% visited with children in their group. The sample was composed of 55% by females and 45% by males. The 83% had a college degree or higher education level (42% college degree and 41% graduate degree).

This data are reported to give a general idea of our sample demographics (i.e. visitors that might have approached and engaged in interaction with Tinker). However, they might not match our sample very well. In fact the experiment was administered in a completely automated fashion and some filtering rules were applied to the data collected. As soon as Tinker detected a visitor approaching the visitor was randomized into one of the three conditions. Kendon's observations of human greetings were primarily focusing on 1-to-1 interactions and this formed the basis for the design of our previous studies which had a single greeting agent and one user approaching it. In the current scenario most visitors arrive in groups of two or more, therefore we had to implement a series of filtering rules to correctly apply our greeting model and obtain "clean" data by discarding from our dataset the cases as shown in Table 4.9.

Filter #	**Description**	**Reason**
1	Approaches including more than one visitor	Violates the theoretical assumptions
2	Single visitors going backwards while approaching Tinker	Violates the theoretical assumptions
3	Approaches started from intermediate points after the *"Distant Salutation"* (T1)	Violates the theoretical assumptions
4	Visitors walking too fast at a speed above the threshold of 1 m/s[5]	Stimuli not applicable
5	Visitors stopping while approaching[6]	Violates the study assumptions
6	Returning visitors recognized by the hand reader	Violates the study assumptions

Table 4.9: The filtering rules applied to our dataset. In order to obtain valid "clean" data ww discarded the cases indicated by the *Description* column. The *Reason* desribes whether theoretical or study design assumptions were violated.

When more visitors were present Tinker started immediately to beckon them as if there were a single person that reached point T4. These cases were filtered out.

---

[5] In earlier studies the approaching speed was 1.11 m/s, but in the pilot testing at the museum we discovered that 1 m/s was the maximum walking speed that allowed Tinker to exhibit the different nonverbal reactions considering its animation capabilities.

[6] A timeout of 8 seconds started after hitting points [T1-T3] of the approach and Tinker started the beckon animation when it expired. The timeout was reset every time a point was reached.

When a single visitor that approached the exhibit and arrived at the *"Encounter"* point T4 didn't start the conversation session for a period longer than 4 minutes Tinker would initiate the interaction. However, we included in our dataset cases when some passersby appeared from behind while a single visitor was involved in an ongoing "clean" approach.

In sum, we filtered out **15.441** cases out of a total of **30.727**[7] visitors that were detected by Tinker during our data collection. Therefore, the total number of visitors across all conditions participating in our study and contributing with clean data is **15.286**. This number includes also visitors that performed *zero* actions and walked away while approaching Tinker (in a *clean* fashion).

4.4.4 Hypotheses

From the results of our previous studies, we predicted that:

- **H1 (Visitor Actions)**: Approaching visitors in the FRIENDLY TINKER group will perform a higher number of actions compared to the HOSTILE TINKER group that, in turn, will complete a higher number of actions compared to the CONTROL group[8];
- **H2 (Session Duration)**: Given that a visitor has performed all the actions to complete the approach and start a conversation session with Tinker. Visitors in the FRIENDLY TINKER sub group that completed earlier actions will engage in longer sessions with the agent compared to the HOSTILE TINKER sub group that, in turn, will have longer sessions compared to the CONTROL sub group.

4.4.5 Data Analysis

Visitor Actions. A Kruskal Wallis test was conducted to compare this ordinal measure between the three groups. The analysis only revealed a tendency toward significance ($\chi^2(2) = 1.58, p = .45$), therefore **H1** was **rejected**. Table 4.10 shows a summary of the *Visitor Actions* outcome. For each of the three experimental conditions, the first column shows the total number of cases in the given condition,

---

[7] N.b. this number includes also passers by that came across the exhibit and quickly moved away.

[8] The CONTROL was expected to rate lowest because nothing was done to *catch the attention* (i.e. engage the visitor in interaction).

then the columns (from left to right) show the percentage of visitors (raw numbers are in parenthesis) that performed the correspondent action described in the heading prior to walking away, except for the first column that only shows how many visitors walked away without performing any action. For instance, if we examine the FRIENDLY TINKER condition, the row indicates that roughly the 90% of the visitors walked away, then the 1.22% only completed the approach and then left, whereas another 1.42% of them went further by also placing the hand in the reader before walking away, and, finally, the 8.06% performed all of the actions by starting a conversation with Tinker.

CONDITION	**TOTAL # VISITORS**	**VISITOR ACTIONS**			
		Walked Away	**Approach Completed**	**Hand in Reader**	**Session Started**
CONTROL	**5085**	89.99% (4576)	1.39% (71)	1.39% (71)	7.21% (367)
HOSTILE TINKER	**5144**	89.46% (4602)	1.34% (69)	1.78% (92)	7.40% (381)
FRIENDLY TINKER	**5057**	89.28% (4515)	1.22% (62)	1.42% (72)	8.06% (408)

Table 4.10: Summary of the Visitor Actions outcome. For each of the three experimental conditions, the first column shows the total number of cases. Then, from left to right, the percentage of visitors that walked away prior to performing any action is shown (in parenthesis the raw number of visitors), and in succession the number of visitors that performed the action described in the heading prior to walking away is reported.

Table 4.11 describes the data collected about the *Visitor Actions* outcome in terms of retention rates. In other words, for each condition we show the percentage of visitors that Tinker was able to retain after every action was performed by them. Thus, for example, in the FRIENDLY TINKER condition the 10.7% of the visitors completed the approach, then 9.5% went a step further by putting the hand in the reader and, finally, 8% started the conversation with the agent.

Session Duration. In order to obtain a normally distributed dataset and fit the ANOVA model's assumptions, since typically duration outcomes appear with right skewness (i.e. all durations are always > 0 and tend to be log-normally distributed), we applied a $log_2 - transformation$ to the original dataset as recommended in [Bland and Altman G., 1996]. We also detected outliers in our dataset corresponding to subjects who had very short session duration. We opted to ex-

RETENTION RATES

CONDITION	TOTAL # VISITORS	VISITOR ACTIONS		
		Approach Completed	Hand in Reader	Session Started
CONTROL	**5085**	10.01%	8.62%	7.23%
HOSTILE TINKER	**5144**	10.54%	9.19%	7.41%
FRIENDLY TINKER	**5057**	10.72%	9.50%	8.08%

Table 4.11: The retention rates for each group. The percentages indicate visitors that Tinker was able to retain in the given condition as they were going further with the number of actions performed prior to walking away.

clude those who had duration inferior to 10 seconds, since this is the time required to exchange at least one dialogue turn with Tinker.

A D'Agostino skewness test [D'Agostino, 1970] ran on the original dataset revealed a highly significant positive skewness of 2.6 ($SE = 14.07$, $p < .001$), whereas the $log_2{-}transformed$ data excluding outliers had a slight (non-significant) negative skewness of -0.1 ($SE = -0.93$, $p = .35$).

Therefore, a one-way between-subjects analysis of variance (ANOVA) was conducted on the $log_2 - transformed$ data to test for ***Session Duration*** differences among the three sub groups that started a conversation with the agent (N = 1138, CONTROL = 365, HOSTILE TINKER = 378, FRIENDLY TINKER = 395).

There was a trend towards a main effect of Tinker's version on the ***Session Duration***, in particular the geometric mean[9] duration in the HOSTILE TINKER group was higher ($GM = 80.50$, 95% $CI\,[74.67, 86.77]$) compared to the FRIENDLY TINKER group ($GM = 77.92$, 95% $CI\,[72.41, 83.88]$), that in turn, had higher geometric mean compared to the CONTROL GROUP ($GM = 73.77$, 95% $CI\,[68.33, 79.62]$). However, this was not significant ($F(2, 1135) = 1.30, p = .27$), therefore **H2** was **rejected**.

---

[9] A back transformation to the original scale of the *Session Duration* measure has been applied to the log-transformed means of the ANOVA test, this results in geometric and not arithmetic means. The confidence intervals (CI) shown are also back transformed.

4.4.6 Discussion

The hypotheses stated for the museum's study were not supported. We didn't observe any substantial difference in the number of actions taken by visitors of the three study groups as predicted by **H1**. In particular, visitors did not perform a higher number of actions in the FRIENDLY TINKER condition compared to the other two as expected. There was only an interesting trend that emerged from our results showing that a higher number of actions (i.e. first completing the approach, then placing the hand in the reader and, finally, starting a conversation with Tinker), were taken by visitors assigned to the FRIENDLY TINKER group (8%) compared to the HOSTILE TINKER group (7.4%) that, in turn, tended to be higher compared to the CONTROL group (7.2%) as originally predicted.

H2 was also not supported, there weren't longer session duration for visitors in the sub group that started the conversation with the FRIENDLY TINKER version, compared to the other sub groups. The emerging trend in this case was toward an unexpected direction. Session duration tended to be higher in both conditions when Tinker was reacting to approaching visitors compared to the CONTROL group. However, there was an interesting tendency for longer session duration in the HOSTILE TINKER sub group compared to the FRIENDLY TINKER one.

These results were inconclusive. We looked into possible reasons behind this outcome and we came up with three broad study design issues:

1. **Hypotheses were false.** Our predictions were based on the assumption that people would be more likely to engage in interaction (i.e. commit to perform more actions towards starting a conversation session with Tinker) and hold for longer conversation sessions with a friendly version of the agent while approaching. Thus, we assumed that first impressions of the agent would impact their choices. While this was demonstrated to be true in the previous two controlled experiments, museum visitors might arrive at the exhibit area with different social expectations given the robotic appearance and artificial nature of the agent. In particular, they might even not expect Tinker to greet them from distance as another human would do in real life. After all, visitors aim to enjoy exhibits and their first impression of the agent might have played a secondary less important role in terms of deciding to engage in interaction and strike up a conversation.
2. **Measurements were wrong.** We might have obtained noisy data, for example we filtered out many potential good subjects from our analysis. Visitors

alone walking too fast (compared to the walking speed suggested by our theoretical framework) or stopping while approaching the exhibit were filtered out. It is possible that those walking faster may have been the most excited ones that wanted to engage in interaction with Tinker. As for the session duration measurement, this might have been affected by the dialogue contents rather than Tinker's greeting behaviors displayed earlier during the approach, so future work is required to analyze the dialogue turns exchanged with the visitor and know more about the reason why visitors dropped the conversation sessions. However, previous experiments with Tinker in the same environment indicate that even with noisy data it is still possible to observe significant outcomes [Bickmore et al., 2011].

3. **Manipulations failed.** It is possible that our manipulations failed and Tinker was not conveying to visitors the impressions we wanted. We were mainly manipulating gaze and smiling behavior. It is possible that the *graphical appearance* of the agent's face and the particular *behavior realization* (i.e. animations) were not supporting our previous findings. Therefore the resulting animations were misinterpreted.

In order to seek further explanations we followed up on the third problem. The nonverbal behavior planning and realization represent a fundamental part of our theoretical framework, therefore we conducted a web survey aimed at testing the difference between FRIENDLY and HOSTILE TINKER versions in terms of friendliness judgments and first impressions that subjects had right after observing the two versions of the agent.

4.4.7 First Impressions of Tinker in a Web Survey

We questioned whether visitors were able to observe any difference between the two versions of Tinker that they had been exposed to. Although we ran a manipulation check to selectively validate the nonverbal behavior exhibited by Tinker (i.e. smiling behavior and different amount of gaze at user) prior to conducting the experiment, we hadn't ensured that judgments of friendliness were correctly formed between the FRIENDLY and HOSTILE TINKER versions.

Stimuli. We designed a within-subjects study where subjects accessed an online survey and observed in a fully randomized order two videos of Tinker, respectively the FRIENDLY and HOSTILE versions, as they were approaching it at the museum, thus the agent was exhibiting the nonverbal reactions associated with each version in an automated fashion.

Measures. After watching each video, subjects were asked to rate their ***Free Impressions*** of the agent in the same format described earlier in Section 4.2.3. Thus we asked them to write the first three adjectives that came to their minds (only one was mandatory). Then we asked them to report their impressions of ***Tinker's Friendliness*** by using a 9 points Likert-scale ranging from *Extremely Hostile* to *Extremely Friendly*. This assessment of friendliness was also done in our behavior interpretation study and described in Section 4.2.3.

Participants. We had 126 subjects participating in this on-line survey (78 males and 48 females). Participants were recruited via public announcements in our university campus and mailing lists. They were representing 6 nationalities[10] with 87% reporting "Icelandic". They were aged 18-60 with 62% in the 21-30 range. All subjects were well educated. Detailed demographics are provided in the appendix C.3.

Quantitative Analysis. Our hypothesis was that the FRIENDLY TINKER version would have received higher ratings of friendliness compared to the HOSTILE TINKER one. We conducted a within-subjects one-way ANOVA on ***Tinker Friendliness*** ratings, however there weren't significant effects (($F(1, 125) = 1.41, p = .24, \eta_p^2 = .011$), thus we **rejected** our hypothesis.

Qualitative Analysis. We have also done a qualitative analysis of the ***Free Impressions*** adjectives following the same methodology described in 4.2.6. We report here the total counts for the identified categories of adjectives having similar meanings in the following format *(<total count in the HOSTILE TINKER condition>, <total count in the FRIENDLY TINKER condition>)*. For the full dataset already grouped in categories see Section C.4 in addendum.

Tinker was mainly judged as "weird, strange, awkward and odd" (29 times vs. 26 times), "nice, polite, likable and friendly" (21 times vs. 21 times) and "clever, assertive and helpful" (12 times vs. 11 times) in both conditions.

HOSTILE TINKER was more often judged as "mechanical, robotic, fake and artificial" (31 times vs. 22 times) and also "lifeless, absent, dull, stiff, flat, cold" (22 times vs. 14 times) compared to FRIENDLY TINKER.

---

[10] We asked participants to select the nation that most represented their cultural identity from a list of all countries in the world as in earlier studies.

On the other hand, the FRIENDLY TINKER version was more often seen as "funny, playful and childish" (7 times vs. 12 times), "silly, stupid and useless" (17 times vs. 22 times), and considered more often "creepy, unnerving and untrustworthy" (5 times vs. 20 times) than the other.

Regarding the graphical appearance of the agent, a few subjects also reported adjectives such as "evil eyes" and "crappy graphics" (for FRIENDLY TINKER) or "big eyes" (for HOSTILE TINKER).

Discussion. This web survey demonstrated that subjects were not able to observe any difference in terms of friendliness between the two videos of Tinker that they observed. Although the HOSTILE TINKER version was judged as more "mechanical, robotic, fake and artificial" than the FRIENDLY TINKER, and we think this might be due to the lack of nonverbal cues expressed by the former (i.e. no smiling and a fewer gazes at the visitor), the two versions of Tinker were judged equally as "nice, polite, likable and friendly". This led us to believe that the mappings between non verbal behavior exhibited by Tinker and the intended communicative functions in terms of expressing friendly/hostile impressions did not work as anticipated based on our previous findings.

Earlier results showed that smiling agents were rated as friendlier, thus more approachable, but from this study we learned that exhibiting this behavior, in such a public setting, becomes trickier. By analyzing subjects' responses in the web survey, we noticed that the smiling behavior of the FRIENDLY TINKER version led subjects to report adjectives such as "evil and big eyes" or "Tinker is over friendly, the agent seems over excited".

Furthermore, behavior realization mattered. It should be considered that nonverbal behavior performed with other behaviors (in our scenario the combination of smiling and gazing behavior) could have either reinforced meanings of those behaviors or contradicted them as pointed out in [Argyle, 1988]. In fact smile is a multi-faceted dynamic expression that can signal much more than "friendly" as we attempted – it can also indicate rapport, amusement, polite disagreement, sarcasm, frustration, pain and more [Ochs et al., 2010, Hoque et al., 2011, Mehu et al., 2008]. Furthermore, considering the visitor's perception of nonverbal cues and subsequent formation of impressions, gender and sex might have impacted the effects of smiling on social judgments [Mehu et al., 2008].

Finally, it is interesting to note that FRIENDLY TINKER was more often judged as "silly, stupid and useless" in the web survey, this might be related to the trends showing visitors having longer session duration with the HOSTILE TINKER. Furthermore, FRIENDLY TINKER was considered “untrustworthy”. This might have important implications in terms of the agent pedagogical and relational role, since visitors might have trusted more the HOSTILE TINKER version when it came to ask information or directions about the museum exhibits and bond with the agent.

4.4.8 Limitations and Future Recommendations

The real world scenario we have chosen represented a great opportunity but came with hard challenges to deal with at the same time. On one hand it allowed us to have a big body of participants despite the restrictive filtering rules, and made it possible to obtain behavioral outcomes in a natural and spontaneous fashion. The experiment was not controlled and visitors were not primed about the study purposes, which is an appealing point when dealing with first impressions. On the other hand, there were uncontrolled aspects regarding the visitors’ and some environmental and technical limitations that we haven’t considered.

In the following sections we summarize some of these limitations and provide a series of recommendations to overcome these limitation in future similar study designs.

Technical and Environmental Limitations

- The anonymous nature of out study prevented us to know more information about the users. In particular, some important attributes, such as their gender or computer literacy, were impractical to obtain but might have impacted their interpersonal judgments of Tinker.
- We couldn’t have the absolute certainty that visitors where observing Tinker’s reactions while approaching. Although we have discarded invalid approach data as explained in Section 4.4.3, we hadn’t tracked visitors’ gaze direction during the approach. Thus, visitors distracted by others during an approach, or simply shifting their gaze away from Tinker, were not being recognized, therefore were part of our dataset.
- Visitors were surrounded by a “noisy” environment. In particular, the study participants might have been distracted during their approaches toward

Tinker by another exhibit that lies at the left side of it as depicted in Figure 4.7, and by the multitude of other visitors passing by the two corridors at Tinker's right side to enter other Computer Place areas. This could be viewed as competing stimuli.

- The space constraints of the particular area where Tinker was exhibited might have altered the validity of the theoretical models chosen (in particular Kendon's greeting model) when applying them to such real life setting. In fact, due to the limited amount of space available and Kinect's sensing capabilities we had to scale down the distances for performing Tinker's greeting communicative functions.
- Further consideration is needed for potential moderating variables that could be hiding Tinker's effects, for example the time of day or the day of week. Those might be a good indicator variable for the level of noise or traffic around the exhibit which could plausibly be a moderator.

Recommendations for Further Similar Studies

The lack of statistical significance of our results and the limitations discussed above led us to formulate a list of adjustments that need to be considered in order to improve the design of future experiments based on the current study:

1. Mapping from communicative functions to behavior: by moving from a controlled experimental setting to a public space we realized that the same mapping from behavior to function didn't work right away. Facets of a single behaviors (e.g. smiling) needs to be further exploited prior to map the intended communicative functions (i.e. friendliness impressions) to behaviors;
2. Behavior realization: in our controlled experiments we used agents with different realization capabilities (i.e. procedural animation) compared to Tinker (pre-made animations). The transition from one system to another is not guaranteed to deliver the same expected results (though we ran a manipulation check as described in the appendix C.2). Therefore the particular realization of behavior represents a problem to consider when planning to apply previous, or other researchers, findings. This has implications for the use of standardized behavior descriptions like BML;

3. Web evaluation survey: the goal of this evaluation was to find possible explanations for the museum's study outcome. We came up with the issues listed in points 1 and 2 as some of the problems that may have undermined the main study, however an evaluation like this could be run prior to beginning the main study;
4. Technical issues: sensing the visitor's proximity was not enough to obtain valid and natural user data for our study. Visitors' distance tracking needs to be accompanied with gaze direction detection in order to ensure that approaching visitors are correctly observing the agent's nonverbal reactions;
5. Environmental issues: given the importance of interpersonal space in our model, the place where the agent is exhibited needs to be carefully planned and tailored to the specific research goals. The ideal setting would allow the exhibit to be at least 10 meters away from the point where a visitor can sight it, thus allowing the agent to accomplish the greeting functions as suggested by Kendon's model adopting the original distances instead of scaling them down. The presence of other installations along the way can distract approaching visitors. It is also recommended to use a less crowded space, given that the focus is on 1-to-1 approaches;
6. Filtering subjects: the filtering rules might have been too restrictive by excluding from our study potentially useful subjects because of theoretical and technical limitations. Therefore it is important, prior to starting another study, to reduce as much as possible the level of filtering. We learned that this can be accomplished by carefully evaluating the applicability of theoretical models into the particular setting chosen and by assigning a top priority to the technical issues that need to be worked out to prevent too much filtering of potentially good subjects (e.g. including in the dataset visitors walking faster than what we assumed as they might be the excited ones that want to engage in interaction with the exhibit).

4.5 Conclusions

The theoretical framework presented in this chapter suggests an approach aimed at operationalizing and managing a user's first impression of a relational agent during a greeting encounter. We focused on users' first impressions of an agent's personality trait (extraversion) and interpersonal attitude (affiliation). They both represent important information that the users might uptake during a greeting encounter with a relational agent. In turn, the user might predict the potential outcome of a relationship with an agent based on these information. An expectation of a positive outcome could directly translate into acceptance of the agent and increased likelihood of subsequent encounters that in turn facilitates the agent's goal of establishing a relationship with the user.

We managed these agent's impressions by manipulating nonverbal immediacy cues, in particular smiling, gazing and proxemics behavior exhibited during the greeting encounter. These are behaviors that frequently recur in such context as suggested by Kendon's greeting and Hall's proxemics theories. Our framework builds on top of this with subtle manipulations of the agent's nonverbal behavior as a function of the interpersonal distance between the user and the agent. This framework takes also into account the user's own personality when interpreting the agent's nonverbal behavior and making relational decisions in terms of how likely and how often they would interact again with the agent.

The behavior interpretation study described in Section 4.2 provided evidence that users form impressions of a greeting agent's extraversion and affiliation in the very first moments of a brief virtual encounter. The setting was a 3D museum entrance shown on a 19″ LCD monitor where users approached the agents with their own avatar in first and third person perspective views. We discovered that smile, gaze and proxemic behavior distinctively contribute to the impression formation of these two characteristics, in particular agent's proxemic behavior significantly influenced judgments of extraversion (personality trait), agents stepping towards the subject's avatar were rated higher in extraversion, whereas gazing and smiling agents were judged as more friendly (affiliation interpersonal attitude). We demonstrated that camera perspective had no influence on user interpretations of an agent's nonverbal behavior given the interaction level limited to just observing the approaches.

The behavior interpretation partially depended on subjects' own personality, as we discovered that only those scoring low in agreeableness interpreted more gazing friendlier compared to less gazing. We have also found that those scoring low in extraversion interpreted more smiling friendlier and a cue of extraversion compared to less smiling. The neuroticism trait did not show any impact in the behavior interpretation.

In the behavior impact study described in Section 4.3 we found that the nonverbal greeting behavior exhibited by life-sized agents had impact on subjects' relational decisions in terms of how likely and how often they would like to spend time with the agents on future virtual guided tours of a 3D virtual museum. Our major finding was that agents' interpersonal attitude had greater impact compared to personality on subjects' decisions, in particular the agents exhibiting high friendliness behavior were more likely to be encountered again and more often than the ones exhibiting low friendliness. Contrary to what we expected, we did not observe any concordance effect between users' own personality and the agents' level of extraversion or friendliness on relational decisions.

In the study at the Museum of Science described in Section 4.4, we expected to find evidence of an increased number of visitors completing approaches toward Tinker and having longer conversations with the friendly version of the agent exhibiting more smiling and gazing during the approach. The results are inconclusive, however, there seemed to be a tendency showing that the versions reacting to approaching visitors (either friendly or hostile) were attracting more visitors, in terms of number of actions performed that led to a conversation with it, compared to the non-reactive one. On the other hand, the trends were showing visitors having longer conversations with the hostile version once the approach was completed. The lack of statistical evidence prevented us to provide definitive conclusions about this study, however we were able to draw from it a list of important considerations that need to be taken into account for future designs of similar experiments in public spaces.

In conclusion, the theoretical framework presented and data gathered in the evaluation studies described can be applied to the practical design and implementation of a computational solution that allows relational agents to automatically select smile, gaze and proxemic cues during the first greeting encounter with users to manage the desired impressions of extraversion and affiliation of them.

In this thesis, the application of such solution will be demonstrated in the domain of virtual learning environments as shown in Chapter 6. We aimed for a plug-in solution for existing fully working agent architectures that would nicely fit into different platforms ranging from online 3D videogame environments (such as the virtual Reykjavik learning environment that will be discussed later) to real world deployments (such as Tinker). Furthermore, we aimed at representing the communicative functions described in our framework, for example the distant salutation phase of the approach, independently from the actual realization choice made by of our component in terms of nonverbal behavior to exhibit.

Many ECA systems [Cassell et al., 1999, Niewiadomski et al., 2009, Cassell et al., 2001, Vilhjálmsson, 2005, Hartholt et al., 2013, van Oijen, 2007] adopted the strategy of using separate representations to specify the communicative functions and supporting multimodal behavior at two levels of abstraction, where the functional level describes the intents of these ECAs, that is what they need to communicate and the behavioral level determines *how* by instantiating the intent as a particular multimodal realization.

This design strategy constitutes the core idea of a common framework for the design and creation of ECAs named SAIBA [Kopp et al., 2006]. SAIBA supports this separation between the representations of function-related versus behavior-related specifications by defining two interface languages named **Function Markup Language** (FML) [Heylen et al., 2008] and **Behavior Markup Language** (BML) [Kopp et al., 2006, Vilhjálmsson et al., 2007a] respectively (the next chapter provides a more detailed description of the SAIBA framework).

We adopted SAIBA as reference architecture for the implementation of our theoretical framework for two main reasons. First, our solution will be compatible with other existing SAIBA compliant ECA systems. Secondly, the existing BML specification already provides a way to specify the relational multimodal behavior of our greeting agents. However, for the generation of those behaviors a more abstract representation of the communicative functions that they accomplish is needed (i.e. FML). For this reason, the following chapter proposes a standard representation of communicative functions, named Function Markup Language (FML) according to SAIBA, that will be used in our computational solution and will support flexibility of integration and abstraction from behavior realization.

"When persons are present to one another they can function not merely as physical instruments but also as communicative ones. This possibility, no less than the physical one, is fateful for everyone concerned and in every society appears to come under strict normative regulation, giving rise to a kind of communication traffic order."

Erving Goffman (1922 – 1982)

5

Representing Communicative Functions: FML Proposal

5.1 Introduction

In Section 2.4 we described the important communicative functions that nonverbal behavior carries out in face to face multimodal communication. This chapter aims at proposing a language to represent communicative functions. This representation should be both human-readable and, at the same time, applicable to an ECA system for describing its communicative functions, thus easy to process by a machine. The proposal for such language is motivated by a key design concept which is the separation between communicative function and behavior (not only nonverbal but also verbal behavior) and described in the following section.

Communicative Function vs. Behavior

Many ECA systems [Cassell et al., 1999, Niewiadomski et al., 2009, Cassell et al., 2001, Vilhjálmsson, 2005] adopted the strategy of using separate interfaces to specify an agent's communicative function and its communicative behavior at two levels of abstraction, where the functional level determines the intent of the agent, that is what it wants to communicate[1] and the behavioral level determines how the agent will communicate by instantiating the intent as a particular multimodal realization.

This separation can be seen as two independent components where one component represents the *mind* of an agent and the other component represents the *body* [van Oijen, 2007]. During a first greeting with the user, for example, the agent's mind decides what function to accomplish (e.g. distant salutation), while the body receives what the mind decides to communicate and renders it at the surface level, according to available communication channels and capabilities of the agent.

This design strategy has several advantages. First of all the agent's mind can produce decisions and intents independently of the body, so for example the same mind can be used for different agent's embodiment (e.g. virtual vs. robotic) or shared across systems. Second, the same communicative function can be delivered with different surface forms (i.e. verbal or nonverbal behavior in case of ECAs) depending on the mental state of the agent or intended impressions that the agent aims to manage on the user. Thus, an agent that wants to manage an impression of a friendly attitude towards the user might express the same function (distant salutation) by using different nonverbal behaviors compared to another agent that wants to appear hostile.

This strategy and the need for sharing and reusing existing working components to speed up the process of getting full conversational systems up and running, led research groups in the ECA community to propose the SAIBA framework.

---

[1] Communicative functions can be also represented as unconscious intents.

The SAIBA Framework

The SAIBA framework (Situation, Agent, Intention, Behavior, Animation)[2] is the result of an international effort to unify a multimodal behavior generation framework for Embodied Conversational Agents [Kopp et al., 2006].

This framework divides the overall behavior generation process into three sub-processes, as depicted in Figure 5.1, each bringing the level of communicative intent closer to actual realization through the agent's embodiment [Vilhjálmsson, 2009]. The interfaces connecting the components are one at the high level, between intent planning and behavior planning, and another interface at the lower level, between behavior planning and behavior realization. They are called **Function Markup Language** (**FML**) [Heylen et al., 2008] and **Behavior Markup Language** (**BML**) [Kopp et al., 2006, Vilhjálmsson et al., 2007a] respectively, and they are designed to be independent of (1) a particular application or domain, (2) independent of the employed graphics and sound player model, (3) and to represent a clear-cut separation between information types (function-related versus process-related specification of behavior) [Kopp et al., 2006].

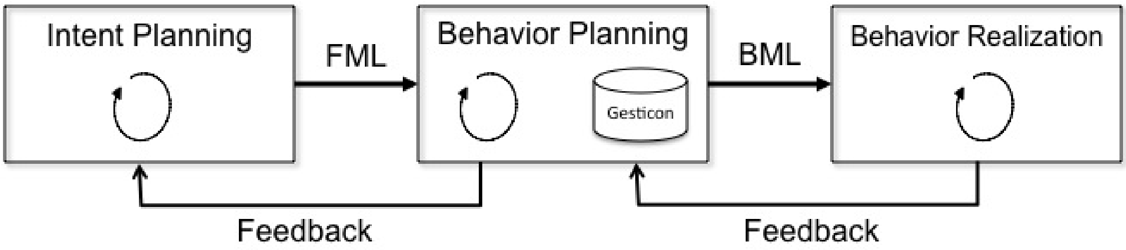

Figure 5.1: The SAIBA framework for multimodal behavior generation, showing how the overall process consists of three sub-processes at different levels of abstraction.

FML describes communicative and expressive functions without any reference to physical behavior, representing in essence what the agent's mind decides. It is meant to provide a semantic description that accounts for the aspects that are relevant and influential in the planning of verbal and nonverbal behavior. An FML description must thus fulfill two tasks. First, it must define the basic semantic units associated with a communicative event. Secondly, it should allow the annotation of these units with properties that further describe communicative function such as expressive, affective, discursive, epistemic, or pragmatic functions [Kopp et al., 2006].

---

[2] Web site: `http://www.mindmakers.org/projects/saiba`.

BML describes the behaviors to express given a function, therefore multimodal behavior should be described so that can be used to control an agent [Vilhjálmsson et al., 2007a]. The last stage handles the realization of the behavior by interpreting the incoming BML and making sure the virtual character behaves accordingly. The behavior realization depends on the particular realization model and can be very diverse. Animations for example can be procedural or fixed and chosen from a repository. Sounds can be generated by a text-to-speech engine or played from file. Therefore what is specified by BML is independent of any specific realization method [Kopp et al., 2006, Vilhjálmsson et al., 2007a].

The framework also presents a *Gesticon*, that can be used by the *Behavior Planner*. This Gesticon is a dictionary which could contain predefined BML behavior definitions.

According to the definition of these two interfaces and the three components in the SAIBA framework, a practical application of our theoretical framework that would be SAIBA compliant can be obtained by decoupling a greeting agent's communicative functions from the exhibited behavior. In particular, at a higher *Intent Planning* level, we want our agent to manage the impressions of personality and attitude towards the user. These high level intents can be decomposed into simpler, lower level, communicative functions that are delivered by the agent over the phases of a greeting encounter as described by our theoretical framework and Kendon's model (e.g. *reaction, distant salutation, close salutation, etc. . .*). The representation of these functions can be done with FML, then at *Behavior Planning* and *Realization* levels these functions are accomplished by generating and exhibiting nonverbal immediacy cues described by the BML language.

Currently BML has been standardized to a first official version adopted by international researchers, however a unified language specification for FML is still work in progress. Furthermore, the SAIBA framework does not specify how the mapping from function to behavior should happen. This mapping is what makes one of our greeting relational agents behave differently from another agent. A different mapping, for example, will result in different realization of communicative functions, thus enabling an agent to employ different impression management strategies.

The next section briefly reviews previous work on FML-like languages and techniques adopted in ECA systems that paved the way of the SAIBA framework. Then, in Section 5.2.3, we discuss the issues that need to be addressed in order to present a unified FML specification that we propose in this thesis (in Section 5.3).

5.2 Towards a unified FML specification

5.2.1 Previous Work

Over the past decade many function and behavior representation languages have been used for the design of ECA systems. This section provides a review of those systems, with particular emphasis on the functional level representation. The behavioral level (i.e. BML-like languages) aspect will not be discussed since it is out of the scope of this thesis. Most of the languages described here are XML based.

Gandalf

In the Ymir architecture, on which the communicative humanoid Gandalf was built [Thórisson, 1997], an *Action Scheduler* (AS) received and executed goals representing both functional and behavioral specifications. High level descriptors of human behavior (user detected in real time) fired goals that were achieved through movement or speech in the Gandalf agent. These higher level descriptors were basic functions to achieve for example: giving-turn, taking-turn, wanting-turn, greet, etc... The AS resolved those functions down to the graphics level and primitive animation commands. Thus, the AS handled, in one place with a single mechanism, both what in SAIBA is referred to as the BML level and FML level.

REA

The separation of intent and behavior was kept in the architecture of the later REA system [Cassell et al., 2000a] by using a fixed messaging pipeline that passed around a multimodal frame. The system was capable of understanding and reacting to natural user input, both audible and visible behavior. A user input *frame* was annotated by an understanding module with two functional interpretations: propositional (content related) and interactional (communicative process related). A decision module used those interpretations for the creation of REA's response in the form of an output frame similar to the input, but this time annotated with the functional descriptions (interactional and propositional) according to the communicative functions that the ECA needed to accomplish. A generation module transformed the output frame into behavior realization that fulfilled the specified functions.

BEAT/SPARK

The Behavior Expression Animation Toolkit (BEAT) was a tool to generate, in an XML based pipeline, multimodal co-verbal behavior based on linguistic and contextual analysis of the text to be spoken [Cassell et al., 2001]. The Spark system modified BEAT for the autonomous generation of avatar multimodal behavior in an online virtual environment based on chat messages exchanged by its users [Vilhjálmsson, 2005]. In Spark the division between communicative function and behavior was made very clear with the definition of two separate XML tag sets. Each message was first annotated with XML tags in terms of various discourse functions related to content and information structure (theme/rheme, emphasis, contrast, topic-shifts) and interaction processes (turn-taking and grounding). These functions were then mapped into supporting nonverbal behavior for a full multimodal delivery by using a separate set of tags. The XML annotation was all done inline with the spoken text and while that made temporal co-occurrence easy to process, it did not allow partially overlapping temporal spans. The term Function Markup Language (FML) was used to describe these tags in the Spark system to contrast them with the set of tags used to describe the supporting Behavior (BML).

MURML

The Multimodal Utterance Representation Markup Language (MURML) [Kranstedt et al., 2002] was a notation system for describing co-verbal gestures produced in synchrony with synthetic speech by the agent Max, a mediator in an immersive 3D virtual environment for simulated assembly and design tasks [Sowa et al., 2001]. A communicative hand gesture was specified by (1) requiring a specific communicative function sufficient for the agent to choose an appropriate behavior from a gesture lexicon, or (2) in terms of the morphological, spatio-temporal features of its meaningful phase (wrist location and trajectory, hand shape, hand orientation), each being described either numerically or symbolically. The main focus of the language was on detailed descriptions of the spatio-temporal features assuming an incremental process model that synthesized continuous speech and gesture in successive chunks marked by internal timing tags. In addition, MURML descriptors allowed augmenting the utterances with facial animations as well as arbitrary body movements.

Social Performance Framework and TLCTS

The Tactical Language and Culture Training System (TLCTS) [Samtani et al., 2008] adopted first an early FML-like proposal described in the Social Performance Framework [Vilhjálmsson and Marsella, 2005] and then aligned with the SAIBA framework. The social performance framework was a precursor of SAIBA and considered intent planning mainly as a decision as to which communicative act to perform. It specified FML in two parts, the first defined certain semantic units associated with the communicative intent (participant, turn, per formative and content) and the second further annotated these units with various communicative functions and expressions (affect, contrast, coping strategy, social goals, meta-cognitive activity, level of certainty, emphasis and illustration).

In the SAIBA compliant TLCTS version the nature of the communicative functions was specified with the following categories of tags: *turn-taking* (take-turn, release turn, keep-turn, assign-turn), *grounding* (initiate, continue, acknowledge, repair, request-repair, request-acknowledgement, cancel), *core speech acts* (inform, wh-question, yes-no question, accept, request, reject, suggest, evaluate, request-permission, offer, promise) and *argumentation* (elaborate summarize, clarify, q&a, convince, find-plan).

Furthermore, they proposed a modulation of those functions (i.e. politeness level) and a unified description of contextual information to apply during the mapping from FML to BML. They proposed an additional third language in the SAIBA framework to represent context that was tentatively called Context Markup Language (CML). With CML, knowledge about the dialog (history of what happened), the environment (e.g. time of the day and current setting) and target culture (culturally appropriate way to express certain functions) is provided in order to be able to generate the appropriate behavior in context.

FML-APML

The FML-APML is an evolution of the Affective Presentation Markup Language (APML) [Mancini and Pelachaud, 2008] developed for the Greta framework [Niewiadomski et al., 2009]. The original APML tags encoded the communicative intentions of an agent following the categorization of [Poggi, 2007]. They expressed the degree of *certainty*, *meta-cognitive* source of information (thinking, remembering, planning), the *speech act* (called performative), *information structure* of the utterance (theme/rheme), *rhetorical relations* such as contradiction or cause-effect (named belief-relations), *turn allocation*, *affect* and *emphasis*.

The set of tags in the FML-APML extended the APML ones with new features regarding the timing and importance of communicative intents, the emotional state of the agent and information on the world. The timing was specified with attributes inspired from the BML recommendations [Kopp et al., 2006, Vilhjálmsson et al., 2007a] and made possible absolute or relative timings of intents with symbolic labels for referencing. This also made it possible to specify communicative functions for non-speaking agents. A common attribute for all tags was the *importance* attribute, depending on its value (ranging from 0 to 1) a communicative functions was encoded differently into behavior to ensure that the meaning was delivered. The emotional state tags gave the possibility to specify an *intensity* (as a numeric parameter from 0 to 1) and a *regulation* type, that was controlling for *felt*, *faked* (emotion aimed at simulating) and *inhibited* emotions (felt but aimed at being inhibited by the agent). The world tag made possible reference to entities in the world and their properties (physical or abstract).

Cultural Influence in Nonverbal Behavior Generation for ECAs

[van Oijen, 2007] shows a SAIBA complaint framework to generate nonverbal behavior given the cultural context and different realization styles. The framework defines the behavior generation in two phases, first the agent's intent is converted into behavior by using a rule based mapping system. Second, styles modulate the generated behavior with additional information about the frequency and expressiveness of certain behaviors.

The communicative functions were expressed with a proposed FML based on both the earlier version of the NVBG module of the Virtual Human Toolkit developed at ICT[3] [Lee and Marsella, 2006] and the model of communicative functions made by [Poggi, 2007]. The FML tags used were mainly divided in two categories: *interactional* and *propositional*. The *interactional* tags combine specifications for the evolution of the discourse structure (speech acts, topic shifts, theme, rheme, etc. . .) with the regulation of the conversation (i.e. turn taking). The *propositional* tags are used to give additional information about the semantic units in the form of sentence parts, phrases and words which already have been structured by using the interactional tags. Such additional information include: emphasis on a word or semantic unit, emotional and cognitive states of the agent, level of certainty about a specific unit, rhetorical relation (i.e. contrast, cause or result), etc. . .

Virtual Human Toolkit

The Virtual Human Toolkit developed at ICT offers a collection of modules, tools, and libraries, as well as a framework and open architecture that integrates these components for the creation of ECAs. It offers coverage of subareas including speech recognition, audio-visual sensing, natural language processing, dialogue management, nonverbal behavior generation and realization, text-to-speech and rendering [Hartholt et al., 2013].

In particular, it uses the *NonVerbal Behavior Generator* module (NVBG) [Lee and Marsella, 2006] to plan the ECA's nonverbal responses by using several FML-inspired concepts, among which are the communicative function (e.g. speech acts and turn management), cognitive operators that drive the gaze state of an agent and elements that relate to emotional states (e.g. affect states such as joy, distress, fear, etc. . .) and coping strategies. The gaze model associates behaviors with what are called cognitive operators by providing a specification of the form and function of gaze patterns. These functions specify detailed reasons behind a particular gaze behavior related to four determinants: conversation regulation, updating of an internal cognitive state (desire, intention...), monitoring for events and goal status and coping strategy. An internal set of rules (written in the XSLT mark-up language) within the NVBG determines which nonverbal behaviors should be generated in a given context.

The nonverbal behavior understanding is kept separated by the functional level, as previously done, with the *MultiSense* module. This module provides the capabilities of both audio-visual sensing and nonverbal behavior understanding, the

---

[3] Institute for Creative Technologies at University of South California.

output messages are broadcast using a new interface language: the Perception Markup Language (PML) [Scherer et al., 2012].

5.2.2 Current Status

The languages for representing communicative functions and the systems adopting those languages listed in the previous section have contributed to a discussion still ongoing about the proposal of a unified FML specification. This discussion has developed over the past years with a series of targeted workshops[4].

Although there were clear differences in the number and kind of dimensions that the systems considered for inclusion in FML, overall there was a big overlap between the various proposals. The work in [Heylen et al., 2008] summarizes the most prominent dimensions, attributes and values associated with them that we consider as the current status for a concrete FML proposal.

As prominent recurring dimensions that emerged (some were discussed in the workshops and not included in the systems reviewed above), there were (1) *contextual information* and *person characteristics* (as participant in a communicative process delivering a function), (2) *communicative actions* (dialogue acts, grounding actions, turn taking), (3) *content* (propositional content, discourse and information structure), *mental and emotional state*, and *interpersonal or socio-relational goals*. The attribute and values proposed for each dimension are briefly summarized as follows (cf. [Heylen et al., 2008]):

Contextual information and person characteristics. *Contextual information* includes cultural and social setting, environmental information (e.g. time of the day), history of interactions and topics discussed. *Person characteristics* are organized in two main dimensions: person information (identifier, name, gender, role, appearance, voice and type as human or agent) and personality.

Communicative actions. The main communicative actions are *turn-taking* actions (e.g. take-turn, want-turn, yield-turn, keep-turn), *grounding* (e.g. initiate, continue, ack, repair, req-repair, req-ack, cancel) and *speech acts* specification (e.g. question, ask, inform, request, etc. . .). However, extra-linguistic or certain non-linguistic actions can also perform certain communicative functions. [Kopp and

---

[4] At Reykjavik in 2005, AAMAS 2008, AAMAS 2009, ICT and Paris in 2010.

Pfeiffer-Leßmann, 2008] suggest the usage of interaction moves referring to a collaborative user-agent task where both dialogue moves are intertwined with manipulative actions.

Propositional content. Assumptions on using formal (logical) languages have been made to represent propositional content. However, there seemed to emerge a first *organization* of *propositions* at *sentence* level (emphasis, given/new information, theme/rheme) and *discourse* level (topics and rhetorical relations between different parts of the discourse).

Mental and emotional states. *Emotional* and *mental states* are believed to contribute to the motivation of a communicative intent. Emotions are divided between *felt*, *faked* and *leaked*. Mental states are defined as cognitive processes such as *planning*, *thinking* or *remembering*.

Socio-relational goals. Social psychological aspects and relational goals play a role in shaping interaction as stated in [Bickmore, 2008]. The concepts of *interpersonal framing* and *relational stance* functions are introduced to affect the behavior produced by an agent with those goals. So, for example, conversational frames might include tags such as *tasks* (information exchange), *social* (for social chat and small talk), *empathy* (for comforting interactions) and *encourage* (for coaching and motivating). Examples of relational stances are *warmth* (high immediacy) or *neutrality* (low immediacy).

The proposed FML dimensions and attributes originated from several factors including the complexity of the systems that suggested them, the different details of specification achieved, the specific demands of the application domains and the theoretical stances taken. Therefore a number of further issues that need to be addressed are still under discussion. We present a summary of these issues in the next section including *current* and several *new* issues that we think are important to face in the specification of FML both for implementing our theoretical framework and, in general, for broadening the FML expressiveness to encompass other SAIBA users' demands.

5.2.3 Further Issues

Defining and separating contextual information

A first question that arises is whether specification of functions (FML) and behaviors (BML) is sufficient for the multi-modal communicative action production. There seems to be the necessity for contextual information, however how much and what information is stored in the context? How is it represented? Do we need a new specific language to represent this information (e.g. CML)?

Person characteristics are an aspect of context that has been considered to be important. In addition to the two dimensions proposed earlier (*person information* and *personality*), are there other dimensions to include? And, is the information included in those two dimensions enough?

The contextual parameters have been shown to be important for the generation of behavior, for example, with respect to the *environmental context* (greetings depending on the time of the day) or *socio-relational* goals. However, we think that in the SAIBA framework a new important issue consists of what is precisely affected by the context not only in terms of behavior generation and realization, but also at the functional level. Does context affect the overall planning of functions? Is one function chosen over another depending on the context?

Defining and classifying functions

A communication function might arise from an action that does not have propositional content, these functions have been classified more generally as *communicative actions*. The main concern is what to consider as a communicative action.

Choosing a classification scheme that embraces all prevailing perspectives on communicative function is not easy, but it will aid the designers of ECAs to use FML at different levels. At a higher level it will be possible to obtain a general outline of the human communicative capacity of a system by noting what general kinds of function specification are available. At a lower level, a designer can expect that functions belonging to the same category will share some specification characteristics and parametrization [Thórisson and Vilhjálmsson, 2009]. The main question that arises is how many groups and categories of functions are needed?

Characterizing and separating conscious vs. unconscious intents

Contextual information does not represent the only determinant for generating communicative functions. It emerged that a broader distinction has to be made between consciously planned intents (for example pure interactional decisions such as the willingness to take the turn) and communicative functions resulting from unconscious determinants such as mental and emotional states. Assuming that a clear characterization can be obtained, the usage of a correct terminology to represent functions that are intended compared to those that are unconsciously produced is required. Furthermore, it remains a question how an unconscious determinant (e.g. emotional state) will affect the production of those communicative functions.

In the FML-APML language, they dealt with the production of unconscious reactions at behavioral level (BML). They argued that for some systems that need to express *mimicry* or *back-channel feedback* behaviors in real time, the intent planning stage of the SAIBA framework should be skipped by providing a short cut from perception of external events (i.e. the user nonverbal and verbal features that require the appropriate mimicry or back-channel feedback) to the immediate realization of behaviors to serve these functions. We think it should still be part of an FML specification the capability of planning immediate reactions, but this raises issues concerning mechanisms to resolve conflicts in cases where an unconscious/reactive intent and a conscious previously planned intent exist at the same time.

Defining temporal constraints and a prioritization scheme

Assigning timing information to a communicative function and allowing coordination among different ones are issues that are certainly important to consider as suggested by [Mancini and Pelachaud, 2008]. [Thórisson and Vilhjálmsson, 2009] suggested that temporal constraints at the functional level of description should be much more coarse-grained than those at the lower levels termed “behaviors” and “execution”. For example, at execution level (BML realization) it is required to deal with frames and milliseconds for the execution of multimodal events, at intermediate level (BML generation) temporal relationships among different multimodal events (such as gaze, posture change, etc. . .) are needed, at functional level it becomes hard to specify exactly how long it will take to accomplish a specific function, nonetheless a designer might still want to be able to specify relative time constraints between functions.

An example of relative timing constraints, suggested by [Thórisson and Vilhjálmsson, 2009] consists of simple synchronization primitives such as *start_together(a,b,...)* or *start_immediately_after(a,b)*, etc. These temporal relationships refer to what they call *"plan chunks"*, therefore *a* and *b* are plan chunks whose relation is described with the primitive, where *b* is the reference. A *"plan chunk"* represents parts of an overall larger and more complex functional plan, for example a set of interrelated propositions to be expressed, for instance how to get from one city to another. Each chunk plan would typically consist of several multimodal acts at the behavioral level.

In addition to allow the specification of temporal constraints, an advantage of cutting FML in *"smaller chunks"* [Bevacqua et al., 2009] or *"chunk plans"* [Thórisson and Vilhjálmsson, 2009] is that they could be processed separately allowing faster generation of corresponding BML compared to a larger FML input (i.e. containing several communicative functions). When dealing with real time reactions (e.g. back channel feedback) a large amount of communicative functions processed as a whole could create an unacceptable delay that would slow down the system's response and make the whole interaction feel unnatural from the user's point of view. However, assuming that FML chunks are adopted in the FML specification. A first issue will be to come up with a precise definition of an FML chunk. Once a definition will be provided, another important issue will be to establish logical boundaries between large sets of communicative functions and identify the subsets that are suitable for being expressed within the same chunk. Finally, in relation to temporal constraints, which are the correct timing primitives that will allow synchronization and coordination among FML chunks?

The decomposition of FML into subsets of chunks also raises two additional issues. One is the need for a feedback system between the intent planner and behavior planner modules of the SAIBA framework as noted by [Bevacqua et al., 2009]. In order to correctly plan and accomplish functions part of an overall communicative intent composed by several FML chunks, the *Intent Planner* module needs to be informed about the current state of the FML that has been transformed into BML (for example some possible states suggested in [Bevacqua et al., 2009] are "playing", "completely played", "discarded" or "interrupted").

The second issue arises from the possibility of having multiple FML chunks that are scheduled at the same time. Conflicting interactions between chunks could be solved by a prioritization scheme, for example in FML-APML [Mancini and Pelachaud, 2008] they proposed the *"importance"* attribute for a function to sort out concurrent communicative functions giving them a priority (i.e. ranging from 0 to 1), but this parameter was also used to choose the multiplicity of multimodal behaviors to accomplish the function. If, for example, importance is low and the plan is about giving directions to reach a particular place, then only some iconic gestures are produced. If importance is very high it adds redundancy by producing more behaviors, for example by looking at the target direction, orienting the body towards it, etc. . .

We think that a prioritization scheme is necessary, in particular assuming the existence of conflicting FML chunks or the co-existence of unconscious/reactive functions and conscious previously planned ones. However, we also think that this prioritization should be kept at functional level merely to establish the sorting of functions to be transformed into behaviors and not the select modalities as proposed with the importance attribute.

Defining an FML document structure

Previous representation languages mainly adopted an XML-like syntax and assigned a nested structure to the specified set of tags (cf. [Mancini and Pelachaud, 2008, Vilhjálmsson and Marsella, 2005]). While it is tempting to adopt a similar structure for a unified FML representation, one may wonder whether this is still a valid solution, and if so, what are the rules that govern the embedding of a set of tags into others. At the current stage of the work, the discussion has been kept on a theoretical level, but this is an important issue when it comes to practically defining a structure for an FML document. In our opinion, there are two stake-holders to consider when defining such structure, on one hand a document should be human readable and easy to understand by a designer, on the other hand it should provide structured information that can be rapidly processed by a machine.

Targeting an FML document to single vs. multiple agents' intents

Another important issue that arises in relation to the structural organization of FML tags is whether an instance of an FML document should refer to a single ECA or multiple ECAs working as part of a unified presentation system. In the former, ECA's communicative functions are represented with a separate dedicated document. Adopting this option allows for distributed processing, therefore might offer better scalability and performance. The latter option requires a central processing module to handle the document and, for example, transform FML into BML. This solution might not be scalable, however, together with temporal constraints it may represent a valid and simple solution to model complex interaction scenarios (e.g. turns management among multiple agents) where functions communicated by an agent need to be planned and coordinated with other agents' functions.

Introducing the concepts of multiple interaction floors and participant's role

In face-to-face communication a person might be engaged in more than a conversation at the same time, and can assume different roles within each one. These considerations provide a departure from the standard dyadic setting towards more complex interaction scenarios. For instance, a person *A* might be engaged in a conversation with *B* and switch to a new conversation with a third person *C* which can be either physically present or not (e.g. on a phone call), but in both cases *C* might not actively participate in the original conversation between *A* and *B*. Another scenario could be a person interacting with another, thus both might assume active speaker/listener roles, and at the same time there might be a third participant that assumes the role of "by-stander" by passively listening or simply being idle. At the FML course in Paris, in 2010, there was common agreement that FML should support the modeling of such scenarios, in particular by introducing the concept of *floor* to model the case of multiple conversations in which a person might be engaged in and including information about who is attending or communicating with whom in each floor and the role assumed within each one.

The idea of modeling more complex interactions involving multiple participants comes with an additional issue. As the number of participants changes, we think this should be reflected, at least, in the generation of behavior (at BML level). Therefore, by enabling a designer to specify at functional level the configuration of interactions, the production of behavior can be affected accordingly. Some examples of configurations might be simple *1-to-1* interactions (for example dyadic), *1-to-many* (for example when describing a public speech) and *many-to-many* (two groups interacting as a whole with each other). In a functional specification of a *1-to-many* interaction for a public speech, for example, we might incorporate this information at functional level as part of the context and the generation of gaze behavior (e.g. the speaker looking at the crowd) might be affected accordingly.

Transforming from FML to BML

The current SAIBA framework does not specify how FML should be transformed into BML. We believe that this aspect cannot be left aside of the framework with individual researchers providing their "*home made*" solutions, as this may critically impact the flexibility of integration into other systems as well as the re-usability of those systems. Therefore, an important issue to face is how this transformation should be achieved. Does it need the adoption of a rule-based mapping system from function(s) to behavior or some different technological solutions? If it will be possible to represent contextual information (for example with CML), how this information will merge with FML in the transformation process?

Applying the FML specification

If we aim at developing a function representation language that will be shared within the ECAs community, there must be core processing components (e.g. for transforming FML to BML as mentioned above) that are made available to and can be used by the community[Heylen et al., 2008]. Furthermore, considering the multitude of SAIBA compliant systems already developed, a challenging issue consists of designing a unified FML representation that will be compatible with those systems and immediately applicable.

5.3 An FML Specification Proposal

In this chapter we present a proposal for a unified FML specification. It builds on top of existing proposed specifications and ideas shown earlier and constitutes a first attempt to address in a single place several of the issues that we have discussed. In light of these issues, we present a proposal of an FML specification based on the categorization of communicative functions described in [Vilhjálmsson, 2009]. It follows that communicative functions are divided in three broad categories identified as **interactional**, **performative** and **mental-emotional state** (a detailed description of these categories will be provided in the remainder of this section). We also aim to represent communicative functions that are either planned intentionally or unintentionally (i.e. unconsciously) by a participant in multimodal interaction. Thus we assume that a communicative function can arise either from a consciously planned communicative intent that the participant aims to accomplish or unconsciously, for example, due to the participant's mental-emotional state. In either case (i.e. intentionally or unintentionally planned) our assumption is that a communicative function represents a goal to achieve in multimodal interaction and based on this assumption we designed our FML representation.

A general overview and the design principles of our proposed FML specification will be shown in Section 5.3.1. A more detailed description will be provided in Sections 5.3.2 and 5.3.3. Finally, some examples demonstrating the application of FML to real case scenarios will be shown in 5.3.4. A description of the full set of FML tags and attributes of this specification can be found in appendix D.

5.3.1 Proposal Overview

Prior to the general overview of our proposal we define several key terms that will appear in the description of the main design principles and in the naming conventions we adopted for tags and attributes:

Participant This refers to an entity (e.g. virtual agent or user) described in a FML document participating in an interaction and, therefore, referred to by the communicative functions as the entity that wants to accomplish it.

Floor A participant can be engaged in several interactions with other participants that we name *floors*. This can be seen as a metaphor for the social contract that binds participants together in the common purpose of interacting.

FML chunk The smallest unit of FML intents associated with a single participant that is ready to be turned into supporting BML-specified behavior.

The following is a summary of the main design principles of our proposal:

- **FML Document structure and target:** We opted for a shared centralized solution, therefore a single FML document instance includes functions that several participants want to accomplish. The language adopted for the representation is XML. A document is divided in two main sections: a *declaration* and a *body*, as illustrated in Figure 5.2. The declarations incorporate contextual information, whereas the body includes all participants' generated functions grouped in FML chunks and belonging to three different tracks as a result of our functions categorization described later.

- **Contextual information and multiple interaction floors:** The declarations section provides a common space to store contextual information, we included these information in the FML specification as opposed to using a separate context markup language. We divided contextual information in two components: a *static* component describing participants information (e.g. gender, age, personality, etc. . .) and a *dynamic* component providing information about the active floors (e.g. participants in each floor).

  Participant information is labeled as *static* since it is meant to endure over time and, therefore, it affects all active floors in which the participant is involved. This specification allows the co-existence of multiple active floors for each participant and each floor involving one or more participants[5].

  The *floors information* component allows the specification of the active floors that will be referenced later in the body and information about the participants in each floor. It has a *dynamic* shape since the information included is meant to be temporarily associated with the particular floor for which it is specified. For instance, a participant's interpersonal attitude can be specified per floor in this component, as we will see later. A detailed description of the information available in each component will also be provided.

---

[5] The existence of a floor with a single participant models an individual that is not interacting with some other participant but is still communicating functions. Therefore, we are assuming that all communicative functions require the existence of a floor.

- **Functions categorization and body tracks:** The body section of an FML document is divided in three sub-sections that we named *"tracks"*. This design reflects the choice of categorizing the communicative functions in three groups, one for each track, as suggested by [Vilhjálmsson, 2009]. The first category of functions (named **interactional**) deals with establishing, maintaining and closing the communication channel, instantiated with a floor, between participants. The second category (named **performative**) covers the actual content that gets exchanged across the communication channel. The third category deals with functions describing mental states and emotions (for simplicity it has been named **mental state**).

  The assignment of functions to different tracks supports the issue concerned with better readability of an FML document from a designer point of view. From the machine point of view, it also aids special treatments for functions in a given category or the embedding of other formal representations targeted to a specific category. For example, it may allow the usage of languages to organize and annotate propositional contents or, as we can see later, to treat mental state functions separately for representing unconscious intents.

- **Temporal constraints and FML chunks:** Splitting the body up into separate tracks requires an overall orchestration of the functions in relation to each other. The order of appearance of functions in the FML document does not necessarily imply anymore the correct delivery time of those. Coarse-grained temporal constraints (described in Section5.3.3) have been introduced to allow the designer a partial ordering of the functions present in a document.

  We also introduced the concept of FML chunks, defined as the smallest unit of functions that are ready to be transformed into BML. The functions included within a chunk are not subject to temporal constraints and can be executed in arbitrary order. The timing constraints allow synchronization and relative timing among chunks across all the tracks.

- **Unconscious intents:** The mental state track assumes a particular meaning that addresses the issue of representing functions that are not deliberately planned by a participant. We took inspiration from the ground state concept in the BML 1.0 standard specification[6]. BML assumes that there is something

---

[6] `http://www.mindmakers.org/projects/bml-1-0/wiki/Wiki`

like a ground state of the ECA. This state comprises several elements, such as the permanent posture or the ground state of the face. For example, when a temporary *posture* behavior ends, the ECA reverts to the posture originally defined in the ground state; when a temporary face expression ends, the face of the ECA reverts to a ground state. In FML, every participant has a *ground state* that comprises his mental and emotional states[7]. Only functions specified in the mental state track can change the participant's ground state for a limited or unlimited time depending on the particular temporal constraint adopted. In essence, the ground state provides additional contextual information about the participant describing internal (emotional or mental) states that can affect the generation and realization of multimodal behavior in the later stages of the SAIBA generation process.

- **Transformation from FML to BML:** As Figure 5.2 shows, both static and dynamic contextual information in the declaration section combined together with functions appearing in the FML body section will affect (1) the production of BML with (2) different realization parameters accordingly, however we are not addressing the issue of generating different functions to accomplish (i.e. generation of FML) based on the contextual information.

The reminder of this chapter describes with more detail the declarations and body sections of an FML document.

---

[7] Other dimensions could be part of the ground state in the future.

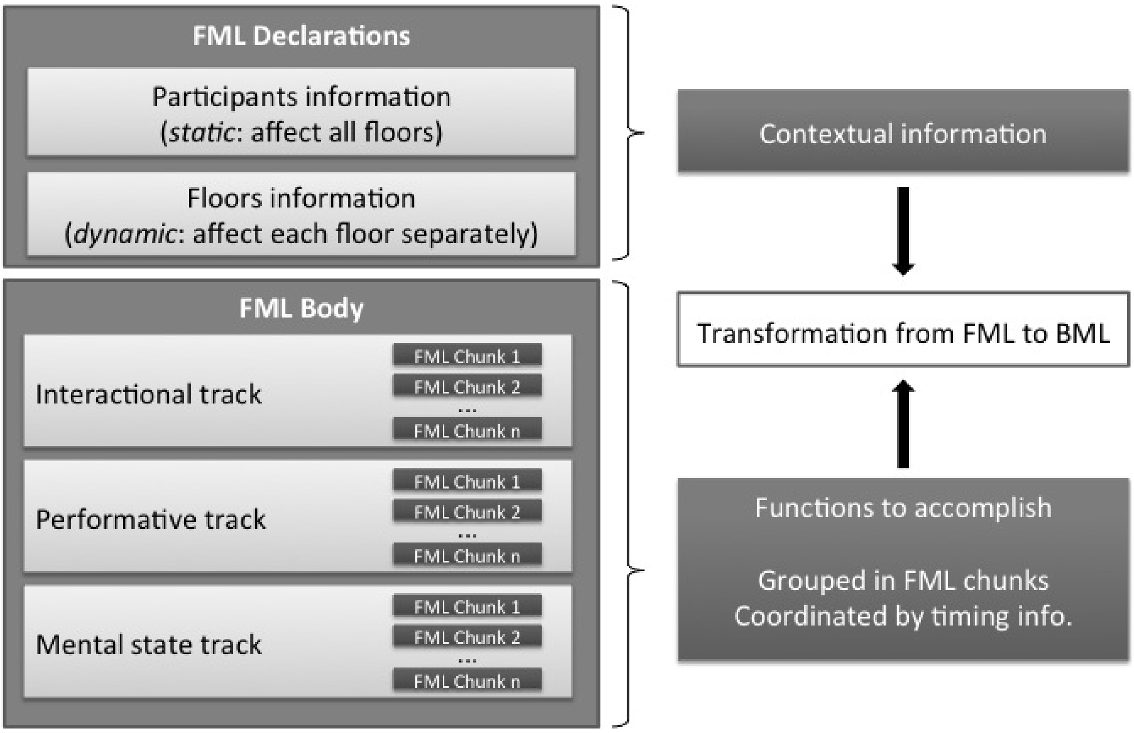

Figure 5.2: An overview of a document adopting our proposed FML specification. A document instance is diveded in declarations and body sections, respectively for contextual information and communicative functions. Contextual information can affect all communication floors (participant information) or slected ones (floors information). The communicative functions are represented within FML chunks that are coordinated among each other with relative temporal constraints.

5.3.2 FML Document: Declarations

The declaration section stores contextual information in two separate sub-sections for participant's information and floors configurations, respectively with two tags named `<identikits>` and `<floors>` as shown in Figure 5.3.

The identikits tag

The `<identikits>` tag contains an `<identikit>` for each participant. The attributes of an `<identikit>` are:

- *id:* A unique identifier associated with the participant. It allows referring to it in an unambiguous way throughout all the other sections of the FML document;
- *name:* A human readable name for the participant;

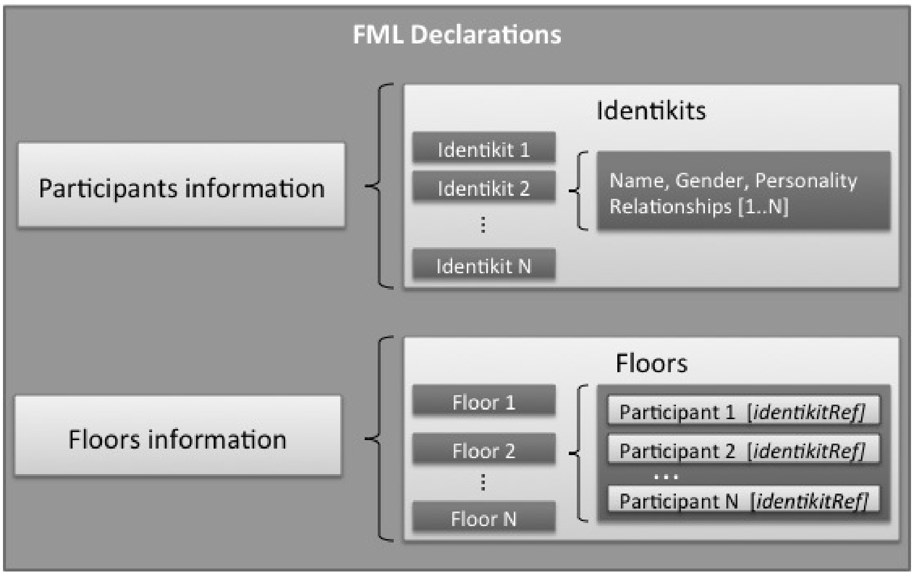

Figure 5.3: The declarations section of an FML document stores contextual information diveded in participants information (identikits) and floors configurations (floors).

- *gender:* The participant's gender[8].

Each **<identikit>** supports the inclusion of the following tags:

- **<personality>**: This tag defines the participant's personality. It has 5 attributes corresponding to each of the Big 5 model dimensions that can have three possible values: *LOW*, *NEUTRAL*, and *HIGH*. For example, an high extroverted and low agreeable participant can be specified as follows:

  ```
  <personality extraversion="HIGH" agreeableness="LOW" />
  ```

- **<relationships>**: This tag allows us to specify the participant's relationship level with other participants. It is a container for one or more **<relationship>** tags. For example, a friend relationship with a participant whose *id* is "Bob" can be specified as follows:

  ```
  <relationship level="FRIEND" with="Bob" />
  ```

---

[8] It is important to note that the contextual information currently defined provide an example of the capabilities offered by our proposal but leaves space for additional information (e.g. cultural background, age, etc. . .) in later versions.

The floors tag

The **`<floors>`** tag contains information about each active floor that the FML document describes. It contains one or more **`<floor>`** tags with the following attributes:

- *floorID:* A unique identifier associated with the floor. It allows functions in the body section to refer to it;
- *floor-cfg:* The configuration of the floor in terms of the type of interaction described. We identified four possible configurations: *individual*, *unicast*, *broadcast* and *multicast*. The naming of these configurations takes inspiration from the network communication protocols terminology. An *individual* configuration describes a single entity (i.e. participant), *unicast* represents the classical dyadic interaction, *broadcast* describes an individual entity interacting with a group and *multicast* characterizes two groups, as a whole, interacting with each other.

A **`<floor>`** can include one or more **`<participant>`** tags depending on the number of participants involved. The **`<participant>`** tag has the following attributes:

- *identikitRef:* A reference to the participant's identikit;
- *entity:* The participant's entity in the given floor configuration as *individual* or *group*;
- *role:* The role assumed by the participant (seen as an entity as specified above) in the given floor. The current proposal draws inspiration from Goffman's participation framework theory [Goffman, 1981] to define this attribute. According to Goffman, participants can have a **speaker** or a **hearer** role, and thereby assume their places in the participation framework for each moment of interaction. Two types of hearers are identified and named in Goffman's participation framework: *ratified* (official) and *unratified* (unofficial) participants. Ratified participants are subdivided into *addressed* and *unaddressed* recipients, and unratified participants or bystanders are subdivided into *eavesdroppers* and *overhearers*, based on their intent and degree of interest. According to this classification, the role attribute can assume the following values: *speaker*, *addressed-hearer*, *unaddressed-hearer*, *eavesdropper* and *overhearer*.

Furthermore, inside a **`<participant>`** tag there can be one or more **`<attitude>`** tags to specify the attitude that the participant has toward another participant in the given floor according to Argyle's status and affiliation model. For example, a floor configured in *unicast* mode (i.e. dyadic) between two participants named "Bob" and "Alice", where Bob has a *friendly* and *submissive* interpersonal attitude toward Alice and both are *individual* entities with *speaker* and *addressed-hearer* roles, can be expressed as follows:

```
<floor floorID="floor1" floor-cfg="unicast">
  <participant identikitRef="Bob" entity="individual" role="speaker" >
    <attitude affiliation="FRIENDLY" status="SUBMISSIVE" towards="Alice" />
  </participant>
  <participant identikitRef="Alice" role="addressed-hearer" entity="individual" />
</floor>
```

FML Listing 5.1: Example of `floor` configuration.

5.3.3 FML Document: Body

The body of an FML document is divided in three tracks following our function categorization scheme. The contents in each track are organized in FML Chunks and timed with relative temporal constraints. Prior to describing the different FML functions that each track can host we will show the specifications for **FML chunks** and **temporal constraints** that can be applied to them.

FML Chunk tag

According to our definition of an FML Chunk provided in the general overview, an `<fml-chunk>` tag has the following attributes:

- *actID:* A unique identifier associated with the chunk;
- *participantRef:* A reference to the participant's identikit *id* attribute.

Figure 5.4 shows how an FML chunk is structured. The first element within a chunk can be a single occurrence of a `<timing>` tag followed by any number of functions defined for the track in which the chunk appears. For example, an FML chunk in the *interactional* track will only include functions categorized as such. The `<timing>` tag allows us to apply a temporal constraint to the whole chunk.

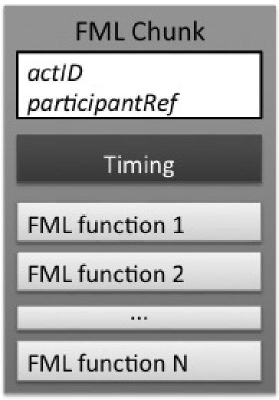

Figure 5.4: The strucure of an FML chunk includes an optional timing tag to assign relative temporal constraints to it and any number of FML functions.

Temporal Constraints

Temporal constraints work on a chunk level with the following design principles:

1. The chunks' order of appearance in the body of an FML document is not meaningful, a scheduler processing the document will take the timing information associated with each chunk and re-order them accordingly;
2. Unless specified by the `<timing>` element, when this information is not present an FML chunk should be scheduled for later processing (i.e. transformation to BML) "as soon as possible";
3. The order of appearance of functions within a chunk is not meaningful and they will be considered in arbitrary order in later stages.

The `<timing>` tag has two attributes:

- *primitive:* Specifies the coarse temporal relationship between the current chunk and the referred one (with *actRef*);
- *actRef:* A reference to another actID (i.e. another chunk). This attribute is required only when the primitive relates to another FML chunk.

The following are possible values for the *primitive* attribute. The names are relatively self-describing, they apply to the chunk in which the primitive is specified and can ask for a reference to another chunk (with the *actRef* attribute). Only "*immediately*" does not require any reference:

- *immediately*
- *must_end_before*
- *execute_anytime_during*
- *start_immediately_after*
- *start_sometime_after*
- *start_together*

The following example shows how temporal constraints can be applied to two FML chunks, *A* and *B*, in the **interactional** track. Chunk *A* has to be scheduled *immediately*, whereas chunk *B* has to be scheduled *immediately after A*.

```
...
<interactional>
  <fml-chunk actID="A" participantRef="part1" >
    <timing primitive="immediately"  />
    <!-- FML Intents -->
     ...
    <!-- FML Intents -->
  </fml-chunk>

  <fml-chunk actID="B" participantRef="part2" >
    <timing primitive="start_immediately_after" actRef="A"  />
    <!-- FML Intents -->
     ...
    <!-- FML Intents -->
  </fml-chunk>
</interactional>
...
```

FML Listing 5.2: Example of temporal constraints applied to a pair of FML chunks.

FML Function Specification

The attributes common to all FML functions tags for each track or category of functions are:

- *floorID:* A reference to the floor in which the communicative function is meant to be accomplished;
- *id:* A unique identifier associated with the function. It allows us to refer to it in an unambiguous way.

Interactional track functions This track supports the specification of a category of communicative functions that serve to coordinate a multimodal interaction. Table 5.1 shows the possible functions that can appear within an FML chunk in this track. The first column on the left side represents a broad category of interactional functions and it is also the name adopted for the correspondent tag. These tags have a common attribute named *type* that narrows down the specification of specific functions within the category. Some of the functions require the specification of the *addressee* attribute. This attribute indicates the participant to which the function is addressed to (functions having this attribute are marked with a "*").

Function Category	Type Attribute
initiation*	*react, recognize, salute-distant, approach-react, salute-close, initiate*
closing*	*break-away, farewell*
turn-taking*	*take, give, keep, request, accept*
speech-act	*inform, ask, request*
grounding	*request-ack, ack, repair, cancel*

Table 5.1: Interactional functions: suggested tag names on the left and possible type attribute values on the right. Functions marked with "*" have an *addressee* attribute that allows to specify to which participant the communicative function is addressed to.

The initiation and closing categories describe the communicative functions, respectively, to manage the initial and termination phases of the interaction. In particular, the different available types of initiation and closing functions are based on the stages of a greeting encounter as suggested by Kendon's greeting model described in Section 2.6. The *approach-react* type for the initiation function has been introduced to specify a custom communicative function reflecting our agent's custom reaction during the approaches, as described in Sections 4.2, 4.3 and 4.4.

The turn-taking, speech-act and grounding functions have type attribute values following the suggestions in [Vilhjálmsson, 2009].

Performative track functions The various functions in this category can be divided across different organizational levels, from the largest organizational structure of a discourse down to the specification of each proposition. In our proposal, the performative track acts as place holder for further embedded extensions of FML specifically targeted to describe performative functions. Therefore, chunks in this track can host one or more `<performative-extension>` tags.

This tag is merely a stub and the description of an extension that will handle its contents is out of the scope of this thesis, though we foresee the inclusion in FML of extension mechanisms similarly to BML[9]. For the current proposal, we limit to show in Table 5.2 a set of possible functions that could be included within a

---

[9] http://www.mindmakers.org/projects/bml-1-0/wiki#Extensions

`<performative-extension>` tag following the recommendations in [Vilhjálmsson, 2009].

There could be different categories of performative functions and, in adherence with other function tags of our proposal, each tag representing a function as a *type* attribute that identifies the particular instance within a category. The `<performative-extension>` tag has also an *addressee* attribute that allows to specify to which participant the included performative act is directed to.

Function Category	Type Attribute
discourse-structure	*topic, segment, ...*
rhetorical-structure	*elaborate, summarize, clarify, contrast, emphasize, ...*
information-structure	*rheme, theme, given, new, ...*
proposition	*any formal notation (e.g. "own(A,B)")*

Table 5.2: Performative functions: suggested tag names on the left and possible type attribute values on the right.

Mental state track functions This track accommodates functions that contribute to visible behavior giving off information without deliberate intent. As a starting point, we proposed functions describing mental states and emotions.

Functions in this track are the only ones capable of changing the *ground state* of a participant. The concept of ground state is kept at abstract level in this proposal, but the idea is that it may affect the manner in which other functions get realized, thus modeling the unconscious side of a participant. We do not specify how the ground state should be modeled and how it should be affecting the behavior generation and realization in later stages. However, we provide several design ideas that might be adopted once the specification will be consolidated and available for the community:

1. Multiple functions can appear within a chunk or across different chunks in the mental state track. A sort of prioritization schema or weighting factor is needed to sort them out and establish the impact that they have on the ground state. Therefore, we proposed a common attribute, only for tags representing functions in this track, named *weightFactor* with values in the interval [0..1].

2. We propose that every function appearing in this track gets sustained by default. This means that unless specified with a temporal constraint (e.g. *must_end_before*), a mental state or emotion changes the ground state permanently. However, reverting to a previous state or voiding the effect of a sustained emotion will be possible by specifying the same function again with the same *weightFactor* as it was before or zeroing it.

3. It should also be possible to specify temporal relationships with other functions. For example we might want to allow an emotion, for example, to be sustained only during the accomplishment of another communicative function in the interactional track (e.g. a distant salutation). The temporal constraints among chunks that we introduced will allow this operation (see example at the end of this section).

Two possible functions that can be included in the chunks appearing in this track are listed in Table 5.3. They describe cognitive processes and emotional states in which a participant could be involved during the interaction. As for other functions the *type* attribute describes the particular instances within the category.

Function Category	Type Attribute
cognitive-process	*remember, infer, decide, idle . . .*
emotion	*anger, disgust, embarrassment, fear, happiness, sadness, surprise, shame . . .*

Table 5.3: Mental and emotional state functions: suggested tag names on the left and possible type attribute values on the right.

We based the specification of the `<emotion>` tag on the FML-APML suggestion [Mancini and Pelachaud, 2008]. Therefore, each `<emotion>` has two attributes that allow to specify the *intensity* and *regulation* of the emotion:

- *intensity:* The intensity of the participant's emotional state in a range from 0 to 1;
- *regulation:* This models a participant's felt or expressed emotional state, it can assume three values:
  - *felt:* A felt emotion;
  - *fake:* An emotion that the participant aims at simulating;
  - *inhibit:* The emotion is felt by the participant but is inhibited as much as possible.

The following example demonstrates the usage of an `<emotion>` tag. It expresses happiness felt by "Bob" while doing a *distant salutation* with "Alice". First, the bond between the *emotion* and the *distant salutation* communicative act is established by using the *"start_together"* timing primitive. Since this emotion becomes sustained by default in Bob's ground state, another chunk (*ACTID03*) is needed to void the effect on subsequent communicative functions. This is done by scheduling another chunk including the same `<emotion>` parameters but having the *weightFactor* lowered to 0.

```
...
<interactional>

  <fml-chunk actID="ACT01" participantRef="Bob" >

    <timing primitive="immediately"  />

    <initiation floorID="floor1" id="id1" type="salute-distant" addressee="Alice" />

  </fml-chunk>

</interactional>

...

<mental-state>

  <fml-chunk actID="ACT02" participantRef="Bob">

    <timing primitive="start_together" actRef="ACT01"/>

    <emotion floorID="floor1" id="id2" type="happiness" regulation="felt"
            intensity="0.8" weightFactor="1.0" />

  </fml-chunk>

  <fml-chunk actID="ACT03" participantRef="Bob">

    <timing primitive="start_immediately_after" actRef="ACT01"/>

    <emotion floorID="floor1" id="id3" type="happiness" regulation="felt"
            weightFactor="0.0" />

  </fml-chunk>

</mental-state>
...
```

FML Listing 5.3: Example showing the usage of an **`<emotion>`** tag in the `mental-state` track.

5.3.4 Full FML Examples for Two Scenarios

In this section we show two full FML examples applied to a real life scenario and a virtual interaction among agents. The real life scenario is about ordering a cheeseburger in a diner and has been subject of discussion in the FML courses held at ICT in USA and in Paris in 2010. The second example is closer to the application domain of this thesis and describes the early stages of a greeting encounter between a user and a greeting virtual agent.

Ordering a Cheeseburger

The following FML document instance describes a two floors interaction happening in a diner among three individuals. The participants are Gilda, Pete and George. Gilda is a customer, Pete is the cashier taking orders and George is the one that makes the cheeseburger.

We assume that Gilda, after having approached the cashier, has already placed her order. Thus, the FML document describes a floor where Pete acknowledges the order just placed by Gilda and another floor where Pete requests George to make a cheeseburger. Gilda acts as by-stander in the second floor between Pete and George.

The following is the heading of the FML document showing the **`<declarations>`** section:

```
<?xml version="1.0" encoding="UTF-8"?>
<tns:saiba-act xmlns:tns="http://cadia.ru.is/FMLSpecification"
    xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
    xsi:schemaLocation="http://cadia.ru.is/FMLSpecification FMLSpecification.xsd ">

<declarations>

  <!-- Participants identikits -->
  <identikits>
    <identikit id="idPete" name="Pete" gender="male">
      <personality extraversion="LOW"/>
      <relationships>
        <relationship level="STRANGER" with="idGilda" />
        <relationship level="FRIEND" with="agentGeorge" />
      </relationships>
    </identikit>
    <identikit id="idGilda" name="Gilda" gender="female" />
    <identikit id="idGeorge" name="George" gender="male" />
  </identikits>

  <!-- Floors configuration -->
  <floors>

    <!-- Floor1 is between Pete and Gilda -->
    <floor floorID="floor1" floor-cfg="unicast">
      <participant identikitRef="idPete" role="speaker" entity="individual" />
      <participant identikitRef="idGilda" role="addressed-hearer" entity="individual
          " />
    </floor>

    <!-- Floor2 is between Pete and George with Gilda as by-stander  -->
    <floor floorID="floor2" floor-cfg="unicast">
      <participant identikitRef="idPete" role="speaker" entity="individual"/>
      <participant identikitRef="idGeorge" role="addressed-hearer" entity="
          individual"/>
      <participant identikitRef="idGilda" role="unaddressed-hearer" entity="
          individual"/>
    </floor>
  </floors>

</declarations>
...
```

FML Listing 5.4: The **`<declarations>`** section of the cheeseburger example.

Contextual information appears in the participants' identikits. In particular, Pete's `personality` is defined as *LOW* for the `extraversion` trait and his `relationships` with other participants are specified. Both Pete and George work in the same place, therefore we assumed that they are friends by setting at "*FRIEND*" Pete's level of relationship with George. The relationship information of the other two participants is left out for clarity, but it is straightforward to specify that Gilda considers both other participants strangers for example.

As for the floors, they both describe a *unicast* configuration. In `floor1` both Pete assumes the role of *speaker* and Gilda has *addressed-hearer* role. In `floor2`, Pete is the *speaker*, George is an *addressed-hearer* while Gilda is an *unaddressed-hearer*. They are all represented as *individual* entities in the two floors described.

The following is the **<body>** section of the FML document:

```
...
<body>

  <!-- Interactional track  -->
  <interactional>

    <fml-chunk actID="ACT01" participantRef="idPete" >
        <timing primitive="must_end_before" actRef="ACT02" />

          <grounding floorID="floor1" id="id1" type="ack" />
    </fml-chunk>

    <fml-chunk actID="ACT02" participantRef="idPete" >
        <timing primitive="start_sometime_after" actRef="ACT01" />

        <turn-taking floorID="floor2" id="id2" type="take" />
        <speech-act floorID="floor2" id="id3" type="request"/>
    </fml-chunk>

    <fml-chunk actID="ACT03" participantRef="idPete"  >

        <timing primitive="start_immediately_after" actRef="ACT04" />

        <turn-taking floorID="floor2" id="id4" type="give" />
    </fml-chunk>

  </interactional>

  <!-- Performative track  -->
  <performative>

      <fml-chunk actID="ACT04" participantRef="idPete" >
        <timing primitive="start_immediately_after" actRef="ACT02"  />

        <performative-extension id="id5" floorID="floor2" addressee="idGeorge">
          <discourse-structure type="topic">
          George make a <rhetorical-structure type="emphasis">cheesburger</rhetorical-structure>
          </discourse-structure>
        </performative-extension>
      </fml-chunk>

  </performative>

  <!-- Mental state track  -->
  <mental-state>

    <fml-chunk actID="ACT05" participantRef="idPete" >
      <timing primitive="start_together" actRef="ACT03"/>

      <emotion floorID="floor2" id="id6" type="anger" regulation="fake" intensity="0.7" weightFactor="1.0" />
    </fml-chunk>

    <fml-chunk actID="ACT06" participantRef="idPete">
      <timing primitive="start_immediately_after" actRef="ACT03"/>

      <emotion floorID="floor2" id="id7" type="anger" regulation="fake" weightFactor="0.0" />
    </fml-chunk>

  </mental-state>

</body>
</tns:saiba-act>
```

FML Listing 5.5: The **<body>** section of the cheeseburger example.

In the example above, Pete *acknowledges* the order just placed by Gilda with a `grounding` function, as can be seen in the chunk at line 7. This *must_end_before* the beginning of the second chunk described at line 13. Within this second chunk, Pete switches to the floor with George, he *takes* the `turn` and performs a *speech act* in the form of a *request*.

Starting *immediately_after*, Pete tells George to make a cheeseburger as described in the `performative` track (see 32). The two chunks in the mental state track accomplish this function with a *fake* emotional state of *anger* (see at line 47). This Pete's emotional state is sustained only for the duration of the performative act, afterwards it gets voided as we can see at line 53.

Finally, *immediately_after* that Pete requests George to make a cheeseburger, Pete *gives* the `turn` away as shown at line 20.

The Very Beginning of a Virtual Greeting Encounter

This second example describes a single floor interaction in a virtual environment between a user and a greeting agent. The participants are named, for simplicity, User and Agent. We assume that the User is approaching the Agent with his own avatar and the Agent is waiting by standing still. The initial distance between the two is such that it allows the Agent to accomplish recognition and distant salutation according to Kendon's greeting model.

The following is the heading of the FML document showing the `<declarations>` section:

```
<?xml version="1.0" encoding="UTF-8"?>
<tns:saiba-act xmlns:tns="http://cadia.ru.is/FMLSpecification"
  xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
  xsi:schemaLocation="http://cadia.ru.is/FMLSpecification FMLSpecification.xsd ">

<declarations>

  <!-- Participants identikits -->
  <identikits>
    <identikit id="idAgent" name="Agent" gender="male" >
      <personality extraversion="HIGH" />
      <relationships>
        <relationship level="ACQUAINTANCE" with="idUser"/>
      </relationships>
    </identikit>
    <identikit id="idUser" name="User" gender="female" />
  </identikits>

  <!-- Floors configuration -->
  <floors>
    <floor floorID="floor1" floor-cfg="unicast">
      <participant identikitRef="idAgent" role="speaker" entity="individual" >
        <attitude affiliation="FRIENDLY" towards="idUser" />
      </participant>
      <participant identikitRef="idUser" role="addressed-hearer" entity="individual"
           />
    </floor>
  </floors>

</declarations>
...
```

FML Listing 5.6: The `<declarations>` section of the virtual greeting encounter example.

According to the participants' identikits the Agent has an *HIGH* level of `extraversion` and its *relationship* level with the User is set to *ACQUAINTANCE*. There is a single *unicast* (dyadic) floor of interaction between the two, where the Agent is the *speaker*[10] and the User is an *addressed-hearer* and, in particular, the Agent has a *FRIENDLY* `attitude` towards the User.

---

[10] The participation in interaction only with communicative actions is still considered being in the speaker role who is performing these actions.

The following is the **<body>** section of the FML document:

```
1  ...
2  <body>
3
4    <!-- Interactional track -->
5    <interactional>
6
7      <fml-chunk actID="ACT01" participantRef="idAgent" >
8        <timing primitive="immediately" />
9
10       <initiation floorID="floor1" id="id1" type="recognize" addressee="idUser" />
11     </fml-chunk>
12
13     <fml-chunk actID="ACT02" participantRef="idAgent" >
14       <timing primitive="start_immediately_after" actRef="ACT01"/>
15
16       <initiation floorID="floor1" id="id2" type="salute-distant"
17                addressee="idUser" />
18     </fml-chunk>
19
20   </interactional>
21
22   <!-- Performative track -->
23   <performative>
24
25   </performative>
26
27   <!-- Mental state track -->
28   <mental-state>
29
30     <fml-chunk actID="ACT03" participantRef="idAgent" >
31       <cognitive-process floorID="floor1" id="id3" type="idle" weightFactor="0.8" />
32     </fml-chunk>
33
34   </mental-state>
35
36 </body>
37 </tns:saiba-act>
```

FML Listing 5.7: The **<body>** section of the virtual greeting encounter example.

In this example the Agent's ground state is initially set as *idle* with a `cognitive-process` function in the `mental-state` track (see line 30) that is scheduled as soon as possible since it doesn't include any `timing` information. The first chunk in the interactional track, then, is immediately scheduled and it represents the Agent's recognition of the User in the approach phase (see line 7). This is followed *immediately_after* by a *distant salutation* (see line 13).

The `performative` track is left empty since there are only communicative acts that do not require the usage of propositional acts.

5.4 Final Considerations and Limitations

In this chapter we have proposed a unified FML specification for the SAIBA framework. The design principles of our specification have roots in the work made by the ECA community over the last decade. We examined a series of languages for representing communicative functions that have been adopted to implement a variety of ECA systems in Section 5.2.1, and we summarized the results of this joint effort and the recommendations for a unified FML specification in Section 5.2.2. In Section 5.2.3 we outlined further issues that an FML specification should address.

We addressed several of these issues:

- We targeted the contents of a **single FML document** instance to *multiple participants* as opposed to multiple documents describing communicative functions of each participant separately.
- We defined a **structure** for an FML document composed by a *declaration* and a *body* section including, respectively, *contextual information* and *communicative functions*.
- We opted to include **contextual information** in the document declaration section as opposed to introduce a new context markup language. These information allows us to specify some *person's characteristics* (including participant's gender and personality), interaction *floors configuration* (including participant's role and interpersonal attitudes towards others) and mental states (with the concept of ground state).
- We introduced the concept of **multiple interaction floors** in which a participant can be simultaneously involved. The contextual information about the configuration of such floors allows to specify the role that each participant has in a given floor in which is involved, the interaction scenario (e.g. dyadic vs. more complex group dynamics) and how the participant is represented (e.g. as individual entity or group).
- We envisioned that the **transformation from FML to BML** will be affected by contextual information in two ways. In particular, the participant's identikits (i.e. static person's characteristics that endure over time) affects all the functions communicated by a participant document wise, whereas the information included in the floors configuration dynamically affect only

the floor in which is defined (e.g. a participant exhibiting an interpersonal attitude towards another in a given floor).

- We adopted a **function classification** and **categorization** schema that has been reflected in the FML document structure by separating the body section in three separate tracks for hosting three main categories of communicative functions identified with: *Interactional*, *Performative* and *Mental State*.
- We introduced the possibility to express communicative functions that do not require the usage of propositional acts (e.g. the interactional functions adopted to express the greeting initiation and closing functions).
- We introduced the concept of **FML chunk** and defined a set of **temporal constraints** that allows to describe relative coarse temporal relationships between chunks.
- We introduced the concept of a participant's **ground state** to characterize the separation between conscious and unconscious intents. Functions belonging to the mental state track might represent unconsciously communicated intents that affect a participant's ground state and this, in turn, could impact the generation and realization of BML behavior in later stages of the SAIBA framework.

Some limitations of this proposal should be also considered. The contextual information needs the inclusion of other important determinants discussed earlier, such as participant's culture and socio-relational goals. We think that the logical separation we have made in the declaration section will easily allow the inclusion of such information, for example culture and age could be part of the identikit, while socio-relational goals can be specified per floor basis. However, a complete ontology of the information that need to be included is needed.

We merely introduced the concept of ground state and we have suggested a simple mechanism to affect this state. However, we still need to specify how the ground state information is stored and where. Furthermore, a mechanism that handles this information in the generation and realization of BML needs to be provided.

This last point is very important both for the purpose of this thesis and, in general for the future of the SAIBA framework. When analyzing the issues to be addressed we underlined that contextual information (or ground state of a participant) can have impact across different stages of the framework. At functional level they can impact the production of functions (i.e. FML), at behavioral level they can impact the generation of multimodal behavior (i.e. BML) and how this behavior is realized (i.e. realization parameters). For this proposal we have chosen to deal with the last two, as we will see in the next chapter, when transforming from FML to BML.

However, there seems to be a demand for inclusion in the SAIBA framework of an external standardized mechanism to handle this transformation and also a specification that goes beyond the mere representation of the two interface languages (FML and BML) is needed. This specification would deal with the contextual information representation and a standardized mechanism for the transformation becomes really important to avoid different researchers adopting their own diverging strategies. We suggest further discussions on this point in future workshops regardless of whether contextual information will be represented inside FML, as we are doing in our proposal, or in a separate markup language. In general, our recommendation is that SAIBA should not only provide standardized interface languages but also techniques and modules that enable to properly transfer between SAIBA components the information represented by these interface languages.

The issue of whether a unified FML specification copes with existing architectures and systems is relatively hard to address. On one hand the design should be taking into account existing architectures as suggested, but on the other hand we think that proposing and consolidating a valid solution for a standard FML specification is an iterative process that needs a starting point and several refinements that can take place only once the community will start to adopt it. Therefore this issue could be faced in further reiterations when the specification will be applied to such existing systems.

We introduced a simple prioritization schema for mental state functions with the *weightFactor* attribute, however an overall prioritization across the three tracks has not been discussed.

Finally, the process of analyzing all the issues to address and the design of this specification led us to some final important considerations.

First, modeling functions and categorizing them, separating and defining contextual information, and in general, dealing with all the aspects of human communicative functions when shaping this proposal required the adoption of a theoretical stance. In this proposal, for example, we adopted specific models of personality (Big 5) and interpersonal attitude (Argyle) to define contextual information. We have also based the categorization of interactional functions on Kendon's greeting model (for example defining the initiation and closing functions and the different types according to the stages of a greeting). This certainly needs agreement among the community, considering also the alternatives (for example other personality models) and the advantages of adopting specific models rather than others. The proposed dimensions in the contextual information and the functions we introduced are just a starting point but this proposal is clearly open for inclusion and feedback from the community.

Secondly, the adoption of temporal constraints requires a mechanism to handle them by properly scheduling the FML chunks that appear in a document. The next chapter will present our solution for this particular problem.

Finally, we think that a system of feedback between the *Intent Planning* module and the *Behavior Planning* module is necessary similarly to the various types of feedback specified between the *Behavior Planner* component and the *BML Realizer* in the BML standard[11].

In conclusion, this proposal addressed several of the issues presented earlier, however it is not yet the ultimate solution. We think that it represents a practical starting point that will spark interesting discussions in subsequent workshops organized by the ECA community. The next chapter will show a practical demonstration that adopts this specification to describe the communicative functions of autonomous greeting agents managing their first impressions in a virtual greeting encounter with their users. The application domain for this practical demonstration is the virtual learning environment for Icelandic language and culture training set in the 3D reconstruction of downtown Reykjavik, in Iceland.

---

[11] `http://www.mindmakers.org/projects/bml-1-0/wiki/Wiki#Feedback`

"Design is not just what it looks like and feels like. Design is how it works."
Steve Jobs (1955 – 2011)

6

Implementing Impression Management for Relational Greeting Agents

6.1 Introduction

This chapter describes the architecture and implementation of a computational solution aimed at (1) providing fully working ECA and RA systems with impressions management capabilities during first greeting encounters with their users, and (2) offering a centralized application domain independent FML processor module to schedule FML and transform it into BML. The solution consists of two software modules that are meant to *enhance* an agent's capabilities with impression management by automating the selection of smile, gaze and proxemic cues, including also body positional and orientational parameters, to convey specific impressions of personality (e.g. extraversion) and attitude (e.g. affiliation) toward the user in the context of a first greeting encounter.

The new FML specification proposed in Chapter 5 and the BML standard represent a possible solution that we adopted for this implementation of our theoretical framework. We wanted to contribute a reusable solution and take advantage of the separation between function and behavior representation, thus SAIBA represented a natural choice in support of these aspects. FML, in particular, is used to represent the greeting communicative functions of an agent as described by Kendon's greeting model in Section 2.6 and our theoretical framework in Chapter 4. The behaviors represented in BML are generated by transforming the FML according to the different impressions of personality and attitude that we are interested in.

Prior to describing the general architecture (in Section 6.3) and implementation (in Sections 6.4 and 6.5) of our modules, we will discuss SAIBA related work (in Section 6.2) that is relevant to some challenges and design issues that we faced along the way.

First, an important design decision was about which technical solution to adopt for the mapping from FML to BML, that is the transformation from one language to another between the SAIBA stages of *Intent Planning* and *Behavior Planning* (the SAIBA framework has been introduced in Section 5.1). Secondly, as part of the implementation, we had to face with the challenge of correctly scheduling FML chunks according to the specified temporal constraints assigned to them.

We limit our review to SAIBA related work. The transformation from FML to BML is a very domain specific problem, therefore we compare other solutions within the SAIBA framework. Regarding the scheduling problem, a comprehensive description of all the work that has been done is too broad and out of the scope of this thesis. In fact, in other application domains the scheduling has been widely afforded (e.g. scheduling problems in the operation research field). We focus on SAIBA related work that has been done in the context of scheduling BML blocks and behaviors. There aren't other works that we are aware of regarding FML scheduling since a standard specification doesn't exist yet.

An application of these modules in the context of Virtual Learning Environments will be shown in Section 6.6. In that section we demonstrate the applicability of our solution to relational greeting agents in the upcoming "*Icelandic Language and Culture Training in Virtual Reykjavik*" project.

Finally, a technical evaluation and some final considerations and limitations of our solution will be discussed, respectively, in Sections 6.7 and 6.8.

6.2 Related SAIBA Work

6.2.1 Transforming from FML to BML

The SAIBA framework specifies two interface languages, FML and BML, to represent respectively communicative functions and multimodal behavior of an ECA. However, it does not provide any specification about how the transformation from FML to BML should happen.

Previous ECA systems have mainly adopted two strategies to solve this problem that are broadly categorized as **data-driven** or **procedural** approaches.

The data-driven approach makes use of *look-up* tables (mainly hard coded in software or with external files often written in XML) describing FML to BML mapping rules or *transformation rules* written in *XSLT*. Look-up tables consists of a set of indexed rules that show possible suggestions to map functions expressed in FML (as index) into multimodal behavior expressed in BML (as result), often a single FML function results in several BML statements that are grouped into BML blocks[1]. XSLT (Extensible Stylesheet Language Transformations) is a template based language for transforming XML documents into other XML documents[2]. XSLT has been used in several ECA systems since both FML and BML are XML-like languages.

In the procedural approach descriptors are transformed from one representation (e.g. FML) to another (e.g. BML) via functions written in a certain programming language. Therefore, the mapping rules are hard coded and the transformation is procedurally accomplished.

The BEAT/SPARK systems [Cassell et al., 2001, Vilhjálmsson, 2005] adopted a mixed approach. There was an XML based pipeline in which FML-like annotated "frames" were transformed into BML-like ones via an XSLT transformation as well as a set of procedures.

The Tactical Language and Culture Training system originally had rules hard coded in the software. Then a user interface was introduced to match communicative functions with context and define behaviors to be produced. This resulted in XML files that stored the mappings in the form of rules to be applied for the transformation from intents to behaviors [Warwick and Vilhjálmsson, 2005].

---

[1] See: `http://www.mindmakers.org/projects/bml-1-0/wiki#Introduction`

[2] XSLT standard: `http://www.w3.org/standards/xml/transformation`

In a later version of the system [Vilhjálmsson et al., 2007b], an agent's reactions at the communicative function level were managed by a *Social Puppet* component coupled with each agent. The social puppet component allowed the agents to accomplish the functions in several ways. First, as programmed procedures for each type of communicative event and as a function of the contextual parameters (e.g. physical configuration, agent standing or sitting, social attitude, hostile or friendly attitude). Secondly, with a direct mapping by using a look-up table, where the indexes were the function and the contextual information and the result was the best matching behavior description or animation name. Finally, the third way to generate behavior was obtained by using an external file containing FML to BML mapping rules that tied up the FML representation with XML contextual representation to produce behavior representation in BML blocks. These files were created by a designer via a graphical user interface as introduced in [Warwick and Vilhjálmsson, 2005].

In a subsequent version [Samtani et al., 2008], the contextual information was represented in a knowledge base language, written in KIF syntax[3], as an ontology of the context (environmental context, previous exchanged contents in dialogue and cultural background). This ontology was used to provide an aid in the translation from FML to BML and was a first proposal for a separate language to describe contextual information (i.e. CML - Context Markup Language).

In [Bevacqua et al., 2009] the concept of FML chunk was introduced and a rule based *FMLtoBML* module accomplished the transformation of chunks into multimodal behaviors expressed in BML.

The cultural framework system described by [van Oijen, 2007] used a *Function Converter* module to transform FML specified functions in BML according to a *Meaning-to-Signal* table storing mapping rules from FML (meaning) to several BML options (signals). The choice among the different options was regulated according to a given probability. This was introduced to provide more variety in expressing behavior.

Finally, the *NVBG* (Nonverbal Behavior Generation) module of the Virtual Human Toolkit [Hartholt et al., 2013] adopts a data-driven approach combining input rules in XML files with XSLT style sheets. The input rules define which words or parts-of-speech trigger which animations targeted to specific cultures, this results in an

---

[3] See: `http://suo.ieee.org/SUO/KIF/suo-kif.html`

intermediate representation that is processed by the NVBG module to generate the final BML to be output.

For the design and implementation of our module we did not explore different XML-based transformation techniques, but reviewing previous SAIBA solutions we decided to adopt a data-driven approach using XSLT since both FML and BML are XML-like languages and XSLT still represents the most diffuse way to transform XML-like documents. It also has the advantage of offering off-the-shelf tools to process the transformation, whereas built-in look up tables require the definition of ad-hoc formats for the files storing the rules and the development of additional software to accomplish the look-up and transformation processes. Furthermore, adopting a data-driven approach allows us to design a centralized module to process FML that is context independent. The mapping rules for our specific application domain (i.e. first greeting encounters) are provided externally via XSLT files, while the processing (i.e. scheduling and transformation from FML to BML) is done independently from the actual context based on the rules provided in input.

On the other hand, it should be said that the programming model and the syntax of XSLT might be unfamiliar and uncomfortable for many procedural-language programmers[4] and it has a steep learning curve that might discourage inexpert designers attempting to create or modify existing style sheets.

6.2.2 Scheduling FML Chunks

The concepts of FML Chunks and temporal constraints that we introduced in our FML proposal (see Section 5.3.3) required us to solve a hard problem. This is the problem of correctly schedule FML Chunks according to the constraints specified and, as a consequence, establishing whether a set of chunks comes with *consistent* timing primitives or not.

For *consistency* we mean that a scheduling plan is feasible, i.e. no circular references are present according to the specific timing primitives chosen. For example, a scenario where for an FML chunk *A* the timing primitive specifies that *A* has to be executed *"immediately_after"* *B* and, in turn, *B* has to be executed *"immediately_after"* *A*, would represent an inconsistent case that has to be detected when attempting to schedule such FML chunks.

---

[4] XSLT is declarative as opposed to stateful and it is based on functional programming ideas.

The only attempt to define a concrete scheduling mechanism for FML that we are aware of is discussed in the FML-APML proposal [Mancini and Pelachaud, 2008]. They proposed to use *start* and *end* timing attributes for FML tags to support the synchronization among them similarly to BML synchronization mechanisms[5]. Therefore, these attributes could assume absolute (numeric non-negative) or relative (to other tags' start and end time attributes) values. The scheduling algorithm to handle FML-APML tags is not discussed, but due to the similarity with BML mechanisms, the solutions adopted for BML scheduling shown in [Reidsma et al., 2011, Thiebaux et al., 2008] represent two valid approaches to consider.

[Reidsma et al., 2011] modeled the scheduling of BML behaviors as a constraint optimization problem. They identified four main constraint types: (1) *explicit* constraints specified as absolute or relative timing values between BML behaviors, (2) *implicit* constraints (e.g. behaviors should have a nonzero duration), (3) *realizer specific* constraints due, for example, to technical limitations of text-to-speech technologies, and (4) *block level* constraints due to the *composition* attribute of a BML behavior[6]. This solution allowed a BML realizer to make on-the-fly adjustments to behaviors while keeping the specified constraints between them intact, therefore without need to rescheduling them all the times and making possible tight mutual behavior coordination between different agents (or agent and user).

In [Thiebaux et al., 2008] (SmartBody system) the behaviors are processed in the order in which they occur in the BML block. The first behavior in the BML block is constrained only by its absolute time constraints and constraint references to external behaviors. Subsequent behaviors are timed so that they adhere to time constraints imposed by already processed behaviors.

Finally, [Zwiers et al., 2011] discuss about chunks of co-expressive speech and gesture behavior that directly map onto BML blocks. They argue that coordination between those two modalities is necessary to obtain natural behavior and that it possible by mapping those modalities into different blocks and introducing synchronization constraints among them. Although the current BML specification may support such synchronization mechanism with *start* and *end* attributes of BML blocks, it might become cumbersome and artificial to specify these constraints when many blocks need to be synchronized. A direct explicit constraint specification might be needed, as we proposed in our FML specification with direct relative timing constraints among FML chunks.

---

[5] See: `http://www.mindmakers.org/projects/bml-1-0/wiki#Synchronisation`

[6] See: `http://www.mindmakers.org/projects/bml-1-0/wiki#Composition`

The usage of *start* and *end* attributes, as suggested by [Mancini and Pelachaud, 2008], entitles a lower level of knowledge of the exact time required to accomplish a function. The approaches described in [Thiebaux et al., 2008, Reidsma et al., 2011] would solve the scheduling problem in this case. However, our goal was to abstract the timing information to a coarse grain level by introducing synchronization primitives that would allow a designer to specify temporal relationships among FML chunks and simply establish a relative order of execution for them. In this way, a designer at *Intent Planning* level does not need to deal with exact duration of those executions compared to later stages in the SAIBA pipeline (i.e. *Behavior Planning* and *Realization*). Therefore, the timing constraints that we proposed in our FML specification are, in principle, similar to the inter-block primitives proposed by [Zwiers et al., 2011] for BML.

As final solution, we modeled the scheduling of FML chunks (as opposed to BML blocks) as a mere **decision problem**. The goal is to decide whether a given schedule for a set of chunks is feasible (i.e. does not contain inconsistent timing primitives among chunks) and then output the chunks in the desired order. The implementation details of this solution are presented in Section 6.5.3, the next section shows a general overview of our greeting modules.

6.3 Design Principles and Architecture

6.3.1 Overview

Figure 6.1 shows an overview of our solution. Each relational agent that aims to manage its impressions during a greeting encounter needs to run a separate instance of a module named `FML Greeting Agent`. The `FML Service` is a single shared module that acts as service provider. It is a general FML processor that keeps track of all the agents in a system, receives concurrent requests as input in the form of FML documents representing communicative functions (adhering to the specification proposed in Chapter 5) and produces BML as output that is properly delivered to the requesting agents in order to be realized. The service module adopts a data driven approach to transform the incoming FML documents into BML. Therefore, it operates independently of the application domain in which it is used. A designer aiming at using it in an application different than ours, can simply specify a different set of transformation rules while using the module that we provide as FML processor. For our specific application domain of first greeting encounters, we focused on mapping rules targeting the *interactional* communicative functions of the *initiation* and *closing* category as described in the FML proposal in Section 5.3.3).

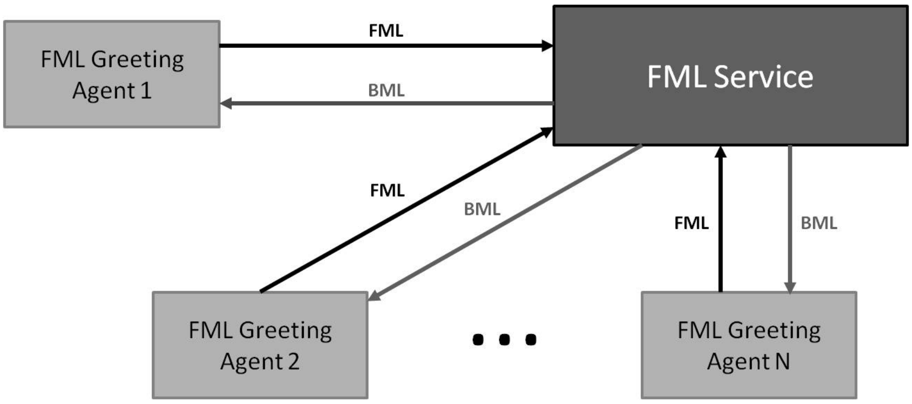

Figure 6.1: An overview of our solution. A separate `FML Greeting Agent` module can be plugged into each relational agent that needs impression management capabilities. A central shared `FML Service` module handles the agents' requests for scheduling and transforming their communicative functions (represented in FML) into multimodal behavior (represented in BML).

In the reminder of this chapter we will adopt the following **naming conventions** for simplicity:

- We will refer to the `FML Greeting Agent` as the **Agent Module** and the `FML Service` as the **Service Module**.
- We will use the term `User` to indicate another entity in a system that begins an approach towards an `FML Greeting Agent`, this entity could be either a user in the real world or another autonomous agent in the 3D world. In practice, for the 3D application of our modules (described later in Section 6.6), with the term *user* we mean a real user represented by an avatar in the 3D environment. However, the `Agent Module` is capable of managing an agent's first impressions towards other approaching autonomous agents as well.

Prior to going into details about these two modules, we summarize here some of the main **design principles**:

- As anticipated, the agent module works on top of existing fully working relational agent systems by enhancing them with impression management capabilities.
- The agent module implements the theories underlined by our theoretical framework described in Section 4. Therefore it handles only 1-to-1 interactions. This means that it can only be used for virtual greeting encounters where a single relational agent is standing still and the *user* approaches it.
- The proposed FML specification supports a shared centralized solution, therefore the service module represents a central processor for FML input documents. According to the specification, a single FML document instance includes functions that multiple participants need to accomplish, thus the service module handles those input documents referring to multiple participants (agents) and enables coordination of their communicative functions throughout the scheduling and BML delivery processes. The variety in the generated BML (e.g. for exhibiting specific behavior to manage different impressions) is obtained via external transformation rules that are given in input to the service module.
- In the practical demonstration of the modules application (described later in Section 6.6), each agent module is attached to a virtual agent in the 3D environment. However, both the FML specification and our implementation support the abstraction from an individual entity (virtual agent) to multiple entities, in order to model group dynamics (for example a group of agents in conversation).

6.3.2 Architecture

The agent and service modules' architecture is depicted in Figure 6.2. In an interactive scenario with several relational agents with impression management capabilities, multiple instances of the agent module are created and each agent gets its own agent module plugged in. The agent module implements an agent's *perception-action* loop with the following components:

Perception This component provides raw measurements taken from the environment via sensors. A sensor that we implemented, for example, detects the raw distance between the agent and the *user*. More sensors could be added to this component later.

Input Understanding It deals with raw data provided by the perception component and turns them into meaningful communicative functions. For example, it informs the reactive planner about the communicative meaning of the raw distances in terms of approach phases described by Kendon's model.

Reactive Planner Given a communicative function (e.g. the current approach phase), this component enables the agent to plan a proper reaction in the form of a communicative function represented with FML.

BML Realizer A SAIBA compliant realizer that receives BML commands and executes them[7].

Init and Error Management These two components are used, respectively, for the initialization and error handling processes.

The communicative functions planned by the `Reactive Planner` in the agent module are transformed into multimodal behavior by the service module. This module acts as a central service provider and includes the following components:

FML Scheduler It detects scheduling conflicts and inconsistencies among the chunks in the input FML document and returns a partial order of the given set of FML chunks by using an internal representation language that will be described later[8].

---

[7] In our implementation we did not have access to a standard BML realizer and we deployed a pseudo-realizer instead that will be discussed later.

[8] This internal representation is a hybrid between FML and BML.

FML-to-BML-Transformer This component transforms the scheduled set of FML chunks into a set of BML blocks that are ready to be delivered to the agents. It uses an external file containing XSLT rules to accomplish the transformation.

BML Dispatcher It dispatches the BML blocks to the correct recipients (agents).

It is important to note that the services provided by these components happen in a sequential fashion. Therefore the service module's interface exposes only the `FML Scheduler` to an agent as entry point, then it runs all the services in chain and outputs the BML.

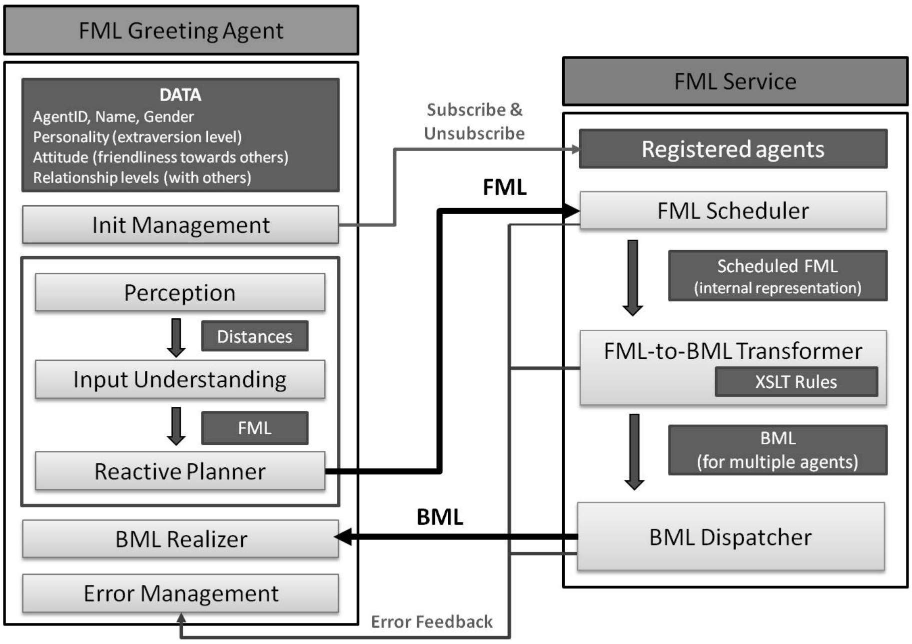

Figure 6.2: The agent and service modules' architecture. The agent module provides an agent with perception, input understanding, reactive planning of greeting communicative functions (in FML) and behavior realization (in BML) capabilities. The service module acts as a central service provider accepting requests from the agents and offering, in chain, FML scheduling, FML-to-BML transformation and BML dispatching services.

6.4 FML Greeting Agent Module Implementation

6.4.1 Setup and Parameters

The setup of the agent module can be done by providing the following parameters (except for the relationship levels):

- **Name**: The name of the agent (doesn't need to be unique, a unique ID is received by the service module after registering to it as described later).
- **Gender**: The gender of the agent, possible values are *MALE*, *FEMALE* or *NOT_SET*.
- **Greeting Awareness**: This is a boolean flag that when turned on (i.e. true) enables the agent impression management during greeting encounters (i.e. activates the perception of an approaching *user* and the execution of reactive behaviors).
- **Extraversion Level (personality)**: The level of extraversion of the agent during greeting encounters. It can assume three values: *HIGH*, *LOW* and *NOT_SET*.
- **Affiliation Levels (attitude)**: the affiliation levels of the agent towards other *users*. This information is stored as pairs of the form `<agentID, level>` where the `agentID` represents an individual agent to which the attitude is referred to during a greeting encounter ("*" can be used to indicate all agents in the system and the real user can be identified with a specific ID that is "player"). The `level` can have the following values: *FRIENDLY*, *HOSTILE* or *NOT_SET*.
- **Relationship levels**: The agent module implements a basic relationship model to keep track of relationship levels with others and vary the behavior produced by an agent after the first greeting encounter depending on the relationship level[9]. This information is stored in pairs of the form <*agentID, relationship level*> where the level can be *STRANGER*, *ACQUAINTANCE* or *FRIEND*. Contrary to all previous data, the relationship levels are not publicly settable, instead are internally handled. Every first time a *user* engages in an encounter the relationship level with it is set to *STRANGER*.

---

[9] See Table 6.3 for a summary of the mapping from greeting functions to behavior including the variations introduced by this basic relationship model.

Upon completing the first greeting encounter (reaching the *Initiation* of conversation phase) the relationship is leveled up to *ACQUAINTANCE*. The second time the *user* completes the encounter the relationship level changes to *FRIEND*. This simple mechanism represents a place-holder that could be easily expanded later with a more detailed model of relationships.

6.4.2 Init Management

Every agent module prior to activating the perception-action loop and in order to be eligible for sending transformation requests to the service module must register with it through a `subscribe` request.

This request is sent by the `Init Management` component. The service module needs this registration to keep track of all the active agent modules in the system. Furthermore, upon a successful registration an agent module receives:

(a) A unique *agentID* used to identify the agent system-wide and fill in the *participant* or *addressee* fields when sending FML documents to the service module.

(b) A set of raw values (in meters) that we identify as *Approach Distances*. These values are used by the `Perception` component of all active agent modules in a system. They are needed when detecting an approaching *user* to establish the phases of a greeting encounter and will be explained later when describing `Perception` and `Input Understanding` components.

(c) *Additional Approach Parameters* that the agent module uses to establish when an approaching *user* is *breaking away* an ongoing approach and when the *farewell* needs to be performed by the agent. The description of these values and their usage will be discussed later when describing the `Perception` and `Input Understanding` components as well.

6.4.3 Error Management

In the current implementation the `Error Management` component only displays a log message when an error arises. Table 6.1 lists all possible errors that the agent and service module can generate. For each error listed an ID, the module that originates it and a description is shown. The *DISPATCHER ERROR* is the only one that is not handled by this component, instead it is logged by the service module directly.

Error ID	Module of Origin	Description
SETUP ERROR	Agent	The `Agent Module` fails to setup due to missing components (e.g. `BML Realizer`) or the registration to the service module fails
PLANNER ERROR	Agent	The `Reactive Planner` fails during the creation of an FML document or placing a transformation request for that document to the service module
REALIZER ERROR	Agent	The `BML Realizer` fails to execute the BML blocks for the agent received by the service module
SCHEDULER ERROR	Service	The input FML document included chunks having inconsistent, conflicting or erroneous temporal constraints assigned to them, thus resulting in an invalid schedule. More details about this problems are provided later in Section 6.5.3
TRANSFORMER ERROR	Service	The `Transformer Service` fails to transform an input FML Document
DISPATCHER ERROR	Service	The BML block to be dispatched refers to an unknown recipient (e.g. an agent that unsubscribed from the service module). In this case the message is logged by the service module directly

Table 6.1: The list of errors that agent and service module can generate. The columns (from left to right) show the error ID, the module of origin and a short description for each error. All the errors are logged by the `Error Management` component of the `Agent Module` except for the *DISPATCHER ERROR*, which is handled by the `BML Dispatcher` component of the `Service Module`.

6.4.4 Perception and Input Understanding

The `Perception` component works in tight cooperation with the `Input Understanding`. The `Perception` provides raw input measurements from the environment via sensors. The current implementation offers a proximity sensor that detects the user's distance from the agent, but later more sensors could be added (e.g. *user*'s speed detection or where the *user* is looking). The `Input Understanding` is a "rule-based" component that applies a set of rules (or conditions) on the raw measurement taken from the perception component to obtain more meaningful functional descriptors. For example, the rules created for our particular context are based on the raw *Approach Distances* values received at registration time. The implemented proximity sensor detects an approaching *user* in real time and depending on these distances several conditions are created. Then, when a condition is met, the input understanding outputs the corresponding functional description of a greeting encounter's phase. Table 6.2 shows the set of rules created and the corresponding descriptor in output when the condition specified by each rule is met. The descriptors are greeting phases corresponding to stages during the approach as described by Kendon's greeting model, except for the *Approach React* phase that has been introduced with our theoretical framework.

Rule (or condition)	**Functional Descriptor**
Distance is **12** meters	React
Distance is less than **9.5** meters	Recognize
Distance is less than **8** meters	Salute Distant
Distance is less than **5.33** meters	Approach React
Distance is less than **3.31** meters	Salute Close
Distance is less than **1.43** meters	Initiate

Table 6.2: A mapping from rules to functional descriptors in our rule-based `Input Understanding` component. A set of rules (or conditions) are turned into meaningful functional descriptors. In the specific example of greeting encounters, rules are based on raw measurements of the *user*'s distance provided by the proximity sensor of the `Perception` component. When a condition is met, the correnspoding functional descriptor describing a *Greeting Phase* of a virtual encounter between a *User* and the agent is output.

These two components operate on two state machines describing (a) the agent states and (b) the phases of a greeting encounter showed, respectively, in Figure 6.3 and 6.4. Central to the perception process is the **React Area**. In terms of raw distances, this is an imaginary sphere centered at the agent's position and having radius as specified by the value (12 meters). In terms of greeting communicative functions it represents the first stage of a greeting encounter. The state machine depicted in Figure 6.3 shows the agent states and the transition from *IDLE* to *REACTING* is based on this area. The agent is initially in the *IDLE* state and it moves to *REACTING* when an approaching user enters the **React Area**. While in the *REACTING* state, the distance perception is activated and a new object named `Approach` is created.

The `Approach` object keeps track of the following information:

- The *addressee* (i.e. the *user*) of the agent's communicative functions and behaviors.
- A unique *floorID* and floor *configuration type* (according to the FML specification) that will be used later to send FML requests to plan the agent's reactive communicative functions.
- The *relationship level* with the *addressee*.
- A list of completed *Approach Phases* to prevent planning the same communicative functions twice during a single approach.
- A flag indicating whether the approach has been completed by the *user* (i.e. the *user* reached the *Initiation* of conversation phase).

While approaching the agent, the *user* can break-away, therefore the agent can go into a state named *SUSPEND REACTING* where the perception is deactivated to allow the agent to accomplish the communicative function associated with this event.

Once the user's approach is completed the agent enters the *DOING CONVERSATION* state but the perception is still active to allow the agent to detect when the user goes away. In the *REACTING* state, when the user keeps moving away from the agent, the latter enters the *DOING FAREWELL* state and exits it when the farewell is accomplished. This transition can happen only once during an active approach.

When the user moves outside of the **React Area** the agent goes again in the *IDLE* state.

The *Additional Approach Parameters* received from the service module allows the agent to handle the break-away and farewell functions. These parameters are the following:

- `Break-Away-Check-Distance` (in meters, default 2m);
- `Break-Away-Check-Time` (in seconds, default 1s);
- Farewell-Check-Time (in seconds, default 2s);

The `Perception` component is able to detect when the user moves away from the agent. During an approach, as depicted in Figure 6.4, there is a transition from any current stage of the approach to the *BREAK AWAY* state when the user moves backwards of more than `Break-Away-Check-Distance` meters during the last `Break-Away-Check-Time` seconds. The farewell is performed when (a) there is an ongoing approach and the agent state is *REACTING TO APPROACH*, (b) the farewell has not been done yet and (c) after `Farewell-Check-Time` seconds that the *BREAK AWAY* state has been entered and the *user* is still moving away from the agent.

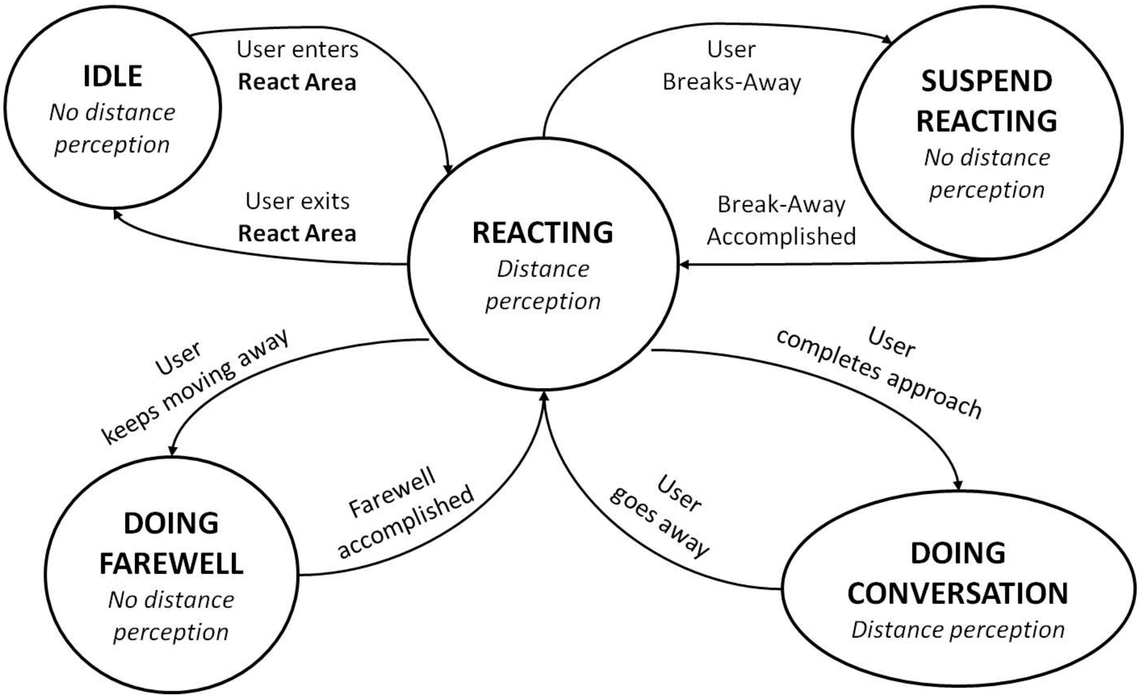

Figure 6.3: The states of an agent controlled by the `FML Greeting Agent` module. The perception of an approaching *user* is active only in selected states and the greeting phases when the agent is in the *REACTING* state are handled by a separate state machine depicted in Figure 6.4.

The state machine depicted in Figure 6.4 shows the phases of a greeting encounter between the *user* and an `FML Greeting Agent`. The initial state is *NO APPROACH* and the transition to *REACT* happens simultaneously with the transition from *IDLE* to *REACTING* in the agent state machine described earlier when the user enters the **React Area**. In this case, however, *REACT* represents a greeting phase as interpreted by the `Input Understanding` component as opposed to the *REACTING* state of the agent.

All other states, in addition to be greeting phases, represent communicative functions that the `Reactive Planner` plans to accomplish as soon as the state is entered. The transitions from one state to another happens according to the *user*'s distance from the agent as detected by the `Perception` component and the corresponding greeting phase indicated by the `Input Understanding` component.

The `Approach` object mentioned earlier is created when the *REACT* state is entered and its flag is marked as completed when the *INITIATION* state is entered. In each state in the diagram shown in the picture, in parenthesis, we show when the communicative function associated with the state is planned only once or multiple times. The list of the *Greeting Phases* already happened during a *user*'s approach it is stored in this object and it simply includes the name of the traversed states in the state machine of Figure 6.4. This information prevents the `Reactive Planner` to plan the same function repeatedly as the user moves back and forth during the approach. We opted to enable only the *INITIATION* function to be planned multiple times since this greeting phase could be reopened repeatedly whilst the *user* is in close proximity of the agent, whereas re-planning other functions, such as close salutation, once happened already seemed odd in our opinion.

The *BREAK AWAY* represents a particular phase of the greeting encounter that, contrary to all the other phases, doesn't happen as the *user* moves closer to the agent but, instead, it happens when the *user* goes away. The condition to establish whether the *user* is breaking away has been described earlier, however, once this state is entered and the respective function is planned, the check after `Farewell-Check-Time` seconds to establish whether the farewell needs to be executed is set.

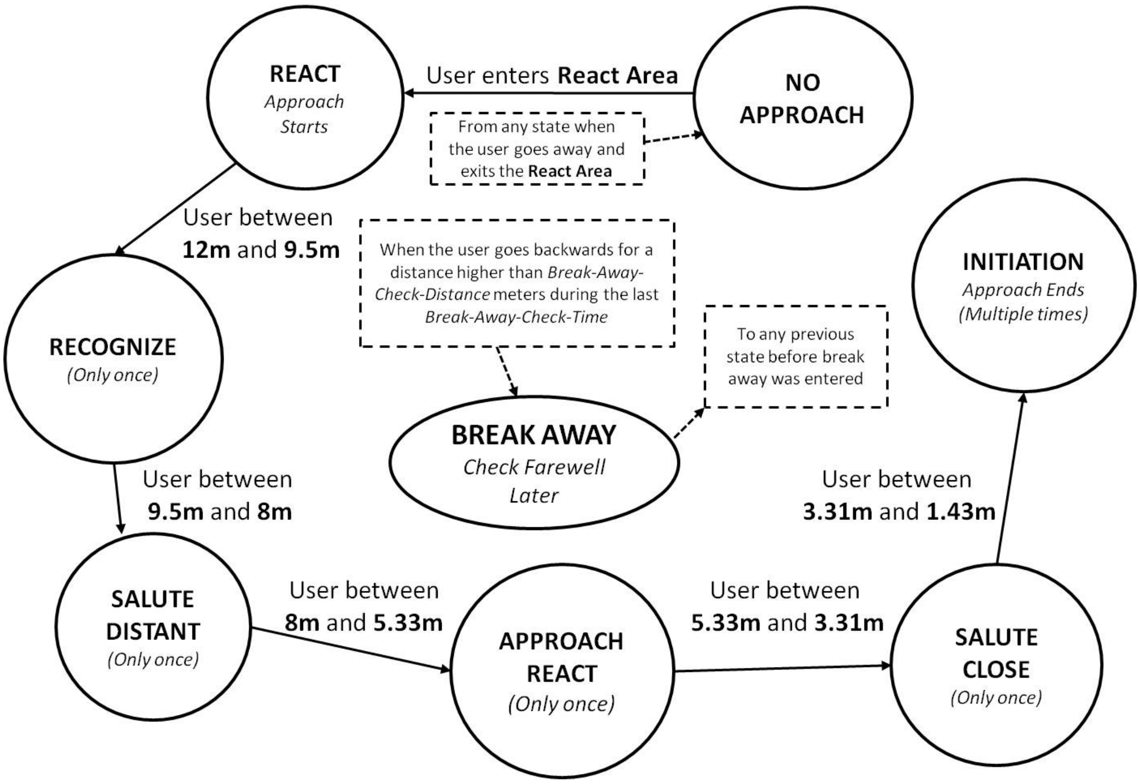

Figure 6.4: The states corrensponding to the phases of a greeting encounter between the *user* and an `FML Greeting Agent`. In each phase the `Reactive Planner` sends a requests to the service module to accomplish the correnspondet greeting communicative function. Some functions, as indicated, are planned only once during a single approach. The *REACT* and *INITIATION* phases represent, respectively, the beginning and ending of an *user*'s approach. In the *BREAK AWAY* state a check to establish if the farewell needs to be accomplished is set.

6.4.5 Reactive Planner

The `Reactive Planner` prepares and sends to the service module FML documents adhering to the standard presented in Chapter 5. The document declaration section is only created when a new `Approach` starts and it is cached for later usage during the greeting phases. The information included in this section are the agent's *personality*, *attitudes* towards others, *relationships levels*, *floorID* and floor *configuration type* (i.e. *unicast* as referring to 1-to-1 greeting encounters).

The body section of a document includes the interactional track functions to accomplish. According to the greeting phase, a reactive communicative function of the "Initiation" or "Closing" category is planned following the mapping of Table 6.3. The table shows a recap of all the raw distances, the correspondent function that is planned and the resulting behavior exhibited by the agent after that the planned FML is sent to the service module and the transformation occurs.

Nonverbal behaviors with the additional description in parenthesis are exhibited by the agent depending on its personality, attitude towards the user or relationship level. When this information is not specified it means that the behavior is always exhibited to accomplish the function indicated. The only gaze behavior without duration information, in correspondence of the *Salute Close* function, indicates a permanent gaze shift of the agent towards the *user*.

Function Category: ***Initiation***

Distance (in meters)	**Function Type**	**Behavior (towards the user)**
12	React	Quick glance (1s)
9.5	Recognize	Gaze (3s) Head toss (only for ACQUAINTANCE) Head nod (only for FRIEND)
8	Salute Distant	Gaze (2s), hand wave and body turn Smile (only when FRIENDLY)
5.33	Approach React	Glance (1s, only when FRIENDLY)
3.31	Salute Close	Gaze, lean forward (only in HIGH extraversion)
1.43	Initiate	Open palm (right hand)

Function Category: ***Closing***

-	Break-Away	Glance around
-	Farewell	Hand wave and gaze away

Table 6.3: The mapping between greeting communicative functions produced by the `Reactive Planner` and the multimodal behavior generated by the service module that is exhibited by the agent with the `BML Realizer` component. The numbers on the left side shows the raw *Approach Distance* values used by the `Perception` component to inform the `Input Understanding` about the *Greeting Phases*.

6.4.6 BML Realizer

The `BML Realizer` component is a SAIBA compliant realizer that receives, asynchronously, BML commands from the `Dispatcher` component of the service module once the `Reactive Planner` has placed a transformation request successfully accomplished.

In our implementation we did not have access to a fully working BML realizer, therefore we implemented a dedicated component named `Pseudo-BML Realizer` that is able to parse and realize only the subset of BML commands that represent the multimodal behavior shown in the rightmost column of Table 6.3.

The agent's facial expressions (i.e. smiling), head movements (nods and tosses), body turn and gaze behavior are exhibited with procedural animation techniques. Whereas the hand waving, open palm and leaning towards the *user* behaviors are performed via standard key frame animations.

Smiling facial expressions are created by controlling Facial Action Units (FAU) [Ekman and Friesen, 1978] affecting the agent's lips, cheeks and eyebrows. The gaze behavior affects both head (at the neck level) and eye movements. Finally, the head toss of our agents is accompanied by an eyebrows raise.

6.5 FML Service Module Implementation

6.5.1 Overview

Each of the three sub-services of the service module run on a separate thread that has a dedicated queue to handle incoming requests with a **First-come, first-served** (FCFS) policy. Figure 6.5 shows the architecture of the service module in detail.

The `FML Scheduler` receives input `Transform Request` objects from the agent modules including the reference to the caller (agent module) and the FML document to transform. These requests are placed in the first empty slot of the *Scheduling Requests* queue. The request is handled by the scheduler and upon successful completion a new document adopting an internal representation language (described later) forms another request for the `FML-to-BML Transformer` service. This request is placed in the *Transformation Requests*. The transformer service, in turn, operates on the request and places the output into the first empty slot of the *Dispatching Requests* of the `BML Dispatcher` service. Finally, the dispatcher handles the request and sends the BML output in the form of BML documents to the appropriate recipient(s).

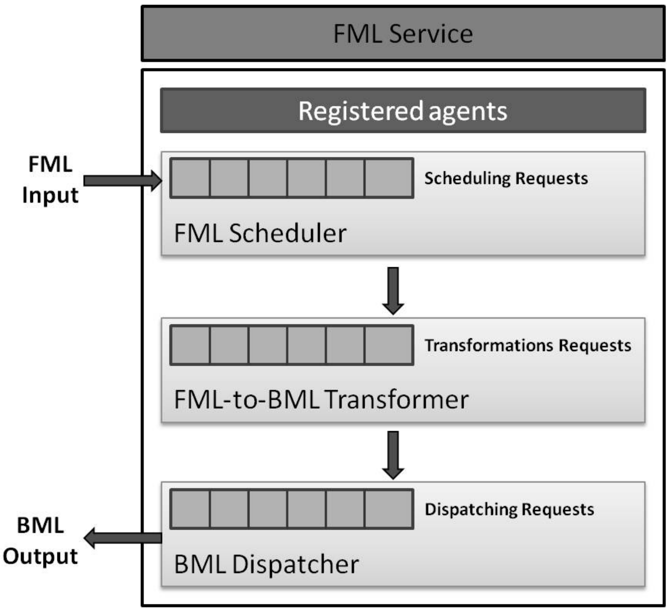

Figure 6.5: A detailed view of the `FML Service` module's architecture. Each sub-service of this module runs on a separate thread and has a dedicated queue to handle incoming requests.

6.5.2 Setup and Parameters

The service module can be set up with the following parameters:

- The mapping between *Approach Distance* values and *Greeting Phases* that are sent to the agent modules after their subscribe requests (Table 6.2 shows the default mapping). This mapping is configurable for future adaptations to other models and theories, however the current defaults values are exactly the same adopted in the studies described in our theoretical framework.
- The *Additional Approach Parameters* described earlier in Section 6.4.2 to handle the break-away and farewell functions.

The *Registered Agents* set stores pairs of the form <`agentID, References to FML Greeting Agent module`> for all the agent modules that are registered (it is internally updated by the service module and empty at start up).

At start up time a thread for each sub-service is launched to start listening for incoming requests.

6.5.3 FML Scheduler

Problem Statement

The `FML Scheduler` component arranges FML chunks such that all temporal constraints are met.

These temporal constraints might require a reference to another FML chunk (e.g. **A** *start_together* **B**, where B is the referred chunk), thus the main challenge is to find out whether a given set of FML chunks comes with consistent temporal constraints that are, for example, not circular or self-referencing and result in a *feasible* schedule of those chunks.

Definition 1 (Feasible Schedule)**.** *A schedule is feasible if it meets the temporal constraints specified in each FML chunk and no circular references among chunks or self-referencing chunks exist.*

As a reminder for the reader, the full set of possible values for the *primitive* attribute of the **`<timing>`** element is listed here:

1. *immediately*
2. *start_immediately_after*
3. *start_sometime_after*
4. *must_end_before*
5. *execute_anytime_during*
6. *start_together*

We will now illustrate with an example how inconsistent temporal constraints need to be detected by the scheduler.

Example 1. Consider a set of three FML Chunks as input with the following **`<timing>`** attributes:

1. Chunk A: `<timing primitive="start_sometime_after" actRef="B" />`
2. Chunk B: `<timing primitive="start_immediately_after" actRef="C" />`
3. Chunk C: `<timing primitive="start_immediately_after" actRef="A" />`

The constraints specified for the three chunks lead to the scenario depicted in Figure 6.6. This scenario represents a schedule that cannot be obtained due to the circular reference preventing **C** from being scheduled *immediately_after* **A**, due to the previous two constraints specified on **A** and **B**.

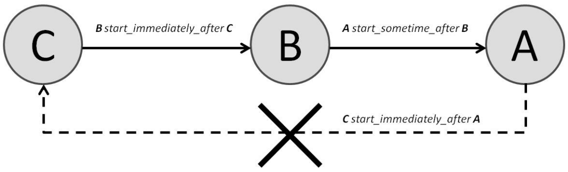

Figure 6.6: An example showing three FML Chunks (A, B and C) represented as circles. The arrows represent the temporal constraints indicated by the primitive of the circle (i.e. chunk) in which they enter. Given the constraints on **A** and **B**, the primitive set for **C** stating that it should *start_immediately_after* **A** leads to an inconsistent schedule.

The remainder of this section includes: A theoretical *formulation* for this problem (see page 159), the *core idea* of our *solution* (see page 161), *further problems* (see page 163) that we had to address due to our specific application domain (i.e. SAIBA framework languages FML and BML), the *design principles* of the scheduling algorithm (see page 167), the description of the *internal representation language* (see page 168) used in the scheduling algorithm's *implementation* (see page 169), *further remarks* about the algorithm (see page 178) and a *final example* (see page 179).

Problem Formulation

We modeled our problem as a mere **decision problem**, therefore given a set of FML chunks as input the goal is to decide if their temporal constraints yield a feasible schedule of the chunks. It is important to note that **we do not provide an optimized schedule** in output since the temporal constraints once assigned earlier by the `Reactive Planner` component (or in general by a designer) are not changeable. Furthermore, the scheduler **does not suggests adjustments** to the given constraints in order to fix inconsistency problems, however it indicates the FML chunks causing these problems.

For this theoretical formulation of the problem we will use the *theory of scheduling* terminology [Brucker, 2007], though we are only dealing with the formulation of a decision problem. This terminology uses the terms **jobs** ("tasks" , "operations" or "activities") and **machines** ("processors", "operators" or "robots"). A **schedule** is for each job an allocation of one or more time intervals to one or more machines.

The corresponding **scheduling problem** is to find a schedule satisfying certain restrictions.

These restrictions are **constraints** (as the temporal ones in our case) that depend on the particular **machine environment** (machines characteristics) and some **job characteristics** (e.g. duration, preemptive vs. non-preemptive, etc. . .). It is out of the scope of this section to fully describe the whole terminology, we recommend the interested reader to consult the introduction in [Brucker, 2007].

We formulate the problem in our specific domain with the following correspondence between a general formulation and the FML specification terminology (including the machine environment and job characteristics):

Jobs The *FML chunks* are the jobs to be executed.

Machines *Participants* are the machines that execute the jobs since each participant accomplishes chunks of functions.

Processing Time The jobs have *unit length* since an FML chunk doesn't specify how long it takes to accomplish the functions included (as discussed at page 101 of Section 5.2.3).

Preemption Jobs are *non-preemptive*, meaning that the processing cannot be interrupted and resumed at a later time.

Precedence Relations There are *precedence relations* between jobs, these are the temporal constraints specified for the chunks.

Dedicated Machine Each job must be processed on a specified *dedicated machine* since each FML chunk refers to a single participant.

Multi-processor Jobs The *start_together* primitive requires that at the same time two jobs need to be executed (e.g. either within the same participant or across different participants, for example when two participants need to accomplish their functions at the same time).

Tight Scheduling Constraint The *immediately_after* primitive introduces a strict requirement on the scheduling. In fact, it doesn't allow any time gap to exist between the previous job (referred by the primitive) and the job that specifies it. This represents a novel challenging issue that is not present in the classical problem formulations since it is very specific to our application domain.

Solution: The Core Idea

The precedence relations between jobs may be represented by a directed acyclic graph (DAG) $G = (V, A)$ where the nodes $V = \{1, ..., n\}$ correspond to jobs, and the arcs to precedence relations. Therefore there exists an arc $(i, k) \in A$ *iff* J_i must be completed before J_k starts. In this case we write $J_i \rightarrow J_k$.

Our decision problem can be solved by building a DAG that represents the given set of jobs and their precedence relations, and by finding a **topological sort** of the graph [Cormen et al., 2009]. If such sort exists the order of the nodes can be returned representing the schedule, otherwise (if a cycle appears) an error arises.

The technique adopted to find a topological sort of DAG consists of visiting G with a **depth first search** (DFS) and order the nodes according to the **reverse farewell time** in the DFS visit. In other words, the descending ordered list of farewell times in the DFS of G gives us the topological sort, if any.

Example 2. To illustrate how this solution works consider a set of 7 jobs $J = \{J_1, ..., J_7\}$ having the characteristics described earlier and, in particular, the following precedence relations, where $J_i \leftrightarrow J_k$ denotes that J_i *starts together* with J_k (and vice-versa) : $J_1 \rightarrow J_2$, $J_2 \rightarrow J_3$, $J_3 \leftrightarrow J_4$, $J_3 \rightarrow J_5$, $J_5 \leftrightarrow J_6$ and $J_5 \rightarrow J_7$.

The DAG representing this set of jobs and their precedence relations is depicted in Figure 6.7. The two numbers on each node indicate the node meeting (left) and farewell (right) time during the DFS visit of the graph. Two nodes at the same depth (e.g. J_3 and J_4) are related with the *start together* constraint. If we relax this constraint for now, the reversed node farewell times gives us the schedule that respects the precedence relations that exist between the nodes.

Therefore the list of farewell times of the nodes after running the DFS on the DAG in descending order:

$$13 \rightarrow 12 \rightarrow 11 \rightarrow 9 \rightarrow 8 \rightarrow 6 \rightarrow 5$$

Corresponds to the schedule:

$$J_1 \rightarrow J_2 \rightarrow J_4 \rightarrow J_3 \rightarrow J_5 \rightarrow J_6 \rightarrow J_7$$

The topological sort of the DAG with reversed DFS farewell time represents the core solution for our decision problem, but as we have seen we relaxed some of the initial constraints. As we can observe from the resulting schedule, by relaxing

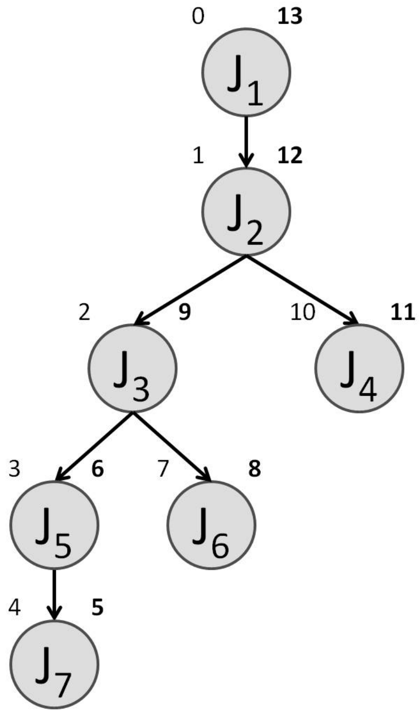

Figure 6.7: A directed acyclic graph representing 7 jobs and their precedence relations. Jobs at the same depth are related by the start together consrtaint. Numbers on the node indicate the meeting (left) and farewell (right) times of the depth first search visit to obtain the graph's topological sort.

the *start together* constraint prior to building the DAG and running the DFS, we miss this information later. In fact we only know that J_3 and J_4 (correctly) start after J_2 but we have no way to specify that J_3 and J_4 must *start together*.

This and several other domain specific problems together with the solutions we found for them will be discussed in the next section.

Further Domain Specific Problems

The two problems presented in this section are originated by the need for incorporating the timing primitives adopted in our FML specification (listed at page 158) in the theoretical solution presented earlier for the construction of a DAG that takes all of them into account.

For two timing primitives we only needed a proper rewording to fit them into our problem formulation while keeping their original semantics. Therefore:

- A *must_end_before* B becomes B *start_sometime_after* A (just equivalent).
- A *execute_anytime_during* B becomes A *start_sometime_after* B (we make sure B starts first).

However, it still remains a problem how to deal with *start_together* and the two primitives indicating the tight *immediately* clause (*immediately* and *start_immediately_after*).

Problem 1: How do we represent *start_together* in our problem formulation? The *start_together* primitive is a constraint that does not belong to the precedence relations often seen in scheduling problems. In order to include this additional constraint, we introduced two new concepts: (a) a special edge type denoted with "$\nrightarrow$" representing an immediate precedence relation and named **Immediate Edge**; (b) A new construct that uses the newly introduced edge and two "fictitious" jobs to represent the *start_together* primitive. This construct is denoted as **Start Together Construct** and is described in the **Example 3**.

Definition 2 (Immediate Edge Notation). *Given two jobs J_i and J_k, the notation $J_i \nrightarrow J_k$ indicates that J_k starts immediately after J_i.*

Example 3. In Example 2, the jobs J_3 and J_4 are supposed to start together. The *Start Together Construct* makes use of the *Immediate Edge* type and two "fictitious" jobs named *Start* and *End* as depicted in Figure 6.8. The jobs J_3 and J_4 are included in the construct and are both scheduled to start *immediately after* the *Start* job, therefore at the same time. The *End* job, in turn, is scheduled immediately after J_3 and J_4 and allows further jobs to keep their precedence relations intact (e.g. $J_3 \rightarrow J_5$ becomes $End \rightarrow J_5$). The introduction of the *Start* node preserves the original temporal constraint between J_3 and J_2. Note that jobs J_5 and J_6 should also be included in a *Start Together Construct* but we left them out for simplicity in the figure.

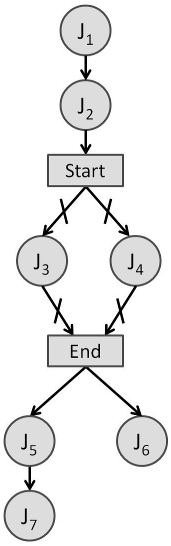

Figure 6.8: The DAG representing the jobs of Example 2 adopting the Start Together Construct to express the temporal constraint that J_3 and J_4 must start together. The construct uses a special edge named *Immediate Edge* and two additional fictitious jobs named *Start* and *End*.

Problem 2: How do we handle the primitives with *immediately* requirement in our problem formulation? The solution for Problem 1 involved the usage of a special edge named *Immediate Edge*. This edge is required also to incorporate the temporal constraint *immediately_after* in the problem formulation. In order to do that, this special edge must receive a specific treatment when building the DAG and running the DFS on it. We have identified three possible cases that our solution must take into account. These cases are described as follows and schematized in Figure 6.9.

CASE 1 (single precedence relation). This is the simplest case. As Figure 6.9 shows, a precedence relation in the form of $J_1 \nrightarrow J_2$ (or J_2 *start_immediately_after* J_1) can be handled by remembering that J_2 has an entering edge of *Immediate Edge* type[10].

CASE 2 (multiple precedence relations all of immediate type). In this case, given a job J_1, there are multiple precedence relations between J_1 and its children but all of the same type, in fact $J_2, J_3 ... J_n$ all require to be scheduled *immediately*

---

[10] In the implementation each node of the DAG will store certain information including the type of the entering edge. The usage of this information by the scheduling algorithm will be described later.

after J_1. The solution for this case, as depicted in Figure 6.9, consists of creating a special *Start Together Construct* that has an *Immediate Edge* entering it (the edge connecting J_1 and *Start*), thus indicating that all children *start_immediately_after* J_1.

CASE 3 (multiple precedence relations of mixed types). This third case is the most generic one, given a a job J_1, there are multiple precedence relations of different type between J_1 and its children $J_2, J_3 ... J_n$. Some of them have normal precedence relations (e.g. jobs $J_3 ... J_n$) and others need to be scheduled *immediately after* (e.g. the job J_2). Figure 6.9 shows the adjacency list of the node representing job J_1 prior to applying our solution (left) and after (right). Adjacency lists are used in the implementation of our algorithm to keep track of all children of a given node . The DFS uses them to visit a node and then all of its children. By placing the nodes with immediately requirement at the rightmost position in these lists we ensure that such nodes will be considered last if the DFS algorithm scans the list of children from left to right. Thus, the last node will be receiving a greater farewell time compared to the previous ones and consequently will appear earlier in the resulting schedule (i.e. ordered by descending farewell times).

Note that the generalization of CASE 3 where more than one children have *Immediate Edge* as a precedence relation, it can still be handled by shifting all children nodes representing jobs that must start *immediately after* J_1 (for example) at the rightmost position in J_1's adjacency list, but then they must be grouped into a special *Start Together Construct* as in case 2. In this way J_1 has a single edge at the rightmost position in its adjacency list pointing to the *Start* node of the construct.

In summary, the solution for CASE 3 requires that the following assumptions are met:

1. The DFS must always begin from the *ROOT* node of the DAG, since starting from random nodes could result in erroneous schedules if, for example, a rightmost node in an adjacency list is picked before others. Thus, starting the visit from the *ROOT* node ensures that the scanning order of the adjacency lists from left to right is always respected.
2. Since (1) must be true, there mustn't exist disconnected nodes that cannot be reached from the *ROOT* node of the DAG[11].

---

[11] Nodes (or jobs) corresponding to FML chunks without timing information could be disconnected from any other, the solution for them will be described in the implementation of the scheduling algorithm.

3. Multiple jobs with *Immediate Edge* precedence relation, when shifted in their parent's adjacency list, must be grouped using the special *Start Together Construct* as underlined in the generalized solution for CASE 3.

4. Only a single temporal constraint can be applied to a job (FML chunk). This avoids multiple entrances to a single node for a job in the DAG. We can make this assumption here since our FML specification conforms to this (see page 115 of Section 5.3.3).

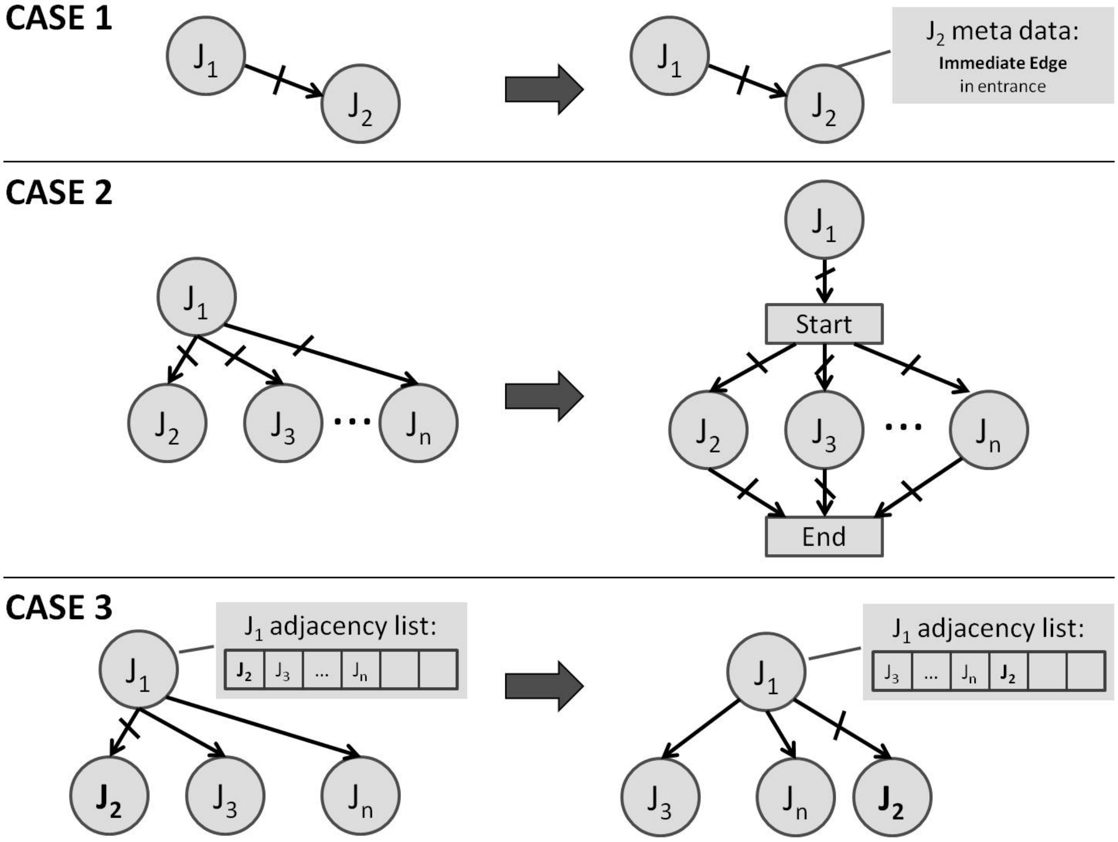

Figure 6.9: Three possible cases that need to be solved in order to incorporate the primitives with *immediately* requirement in our problem formulation. In CASE 1 a single precedence relation of this type is handled. CASE 2 describes the situation where multiple precedence relations all of immediate type exist. In CASE 3 there are multiple precedence relations of mixed (regular and immediate) types.

Scheduling Algorithm: Design

So far we formulated our decision problem, showed the core idea of how to solve it and provided solutions for further domain specific problems regarding the inclusion of all timing primitives in our problem formulation.

We now list the design principles of our scheduling algorithm:

1. The algorithm performs two main operations given the input FML document. First, it builds the DAG with nodes representing the FML chunks in the document. Secondly, in addition to deciding whether the temporal constraints of the FML chunks constitute a feasible schedule, it also prepares the document for making the transformation process as smooth as possible. The `<declarations>` section of the document is kept as it is received by the agent module since it contains information used later by the `FML-to-BML Transformer`. However, the scheduler operates on the `<body>` section adopting elements and attributes of an *Internal Representation Language* that is a hybrid between FML and BML. This language will be briefly described in the following section, its main goal is to facilitate the transformation process and provide information to the transformer service, such as suggestions for BML attributes to use, so that the transformer can simply apply XSLT rules to FML functions without bothering with the order of appearance of these functions in the document and other technical details relative to BML attributes that need to be set.

2. The scheduling of FML chunks in the input document is not order dependent (i.e. order of appearance) but depends on the temporal constraints specified for each chunk.

3. FML Functions included in a chunk are later transformed into BML behavior and executed at arbitrary time (i.e. the functions' order of appearance within a chunk does not matter).

4. The scheduler cannot change or fix temporal constraints specified in the input document by the *Reactive Planner* component of the agent module.

5. The scheduler while preparing the FML document for the transformation modifies the `<body>` structure. The tracks are removed from the original document and the result of the algorithm is a document written in a hybrid internal language including a set of BML blocks, but containing FML functions ready to be transformed instead of BML behavior tags. The order of

these BML blocks in the document matters in this case. As a general principle, the scheduler attempts to put as many functions as possible in the same BML block for a given participant (or character in the BML terminology).

6. FML chunks with *"immediately"* timing primitive assigned to them are grouped by participant and they are scheduled to be executed prior to all other chunks in the document. Grouping by participant means that functions of FML chunks with immediately primitive of the same participant are all "collapsed" into a single BML block.
7. FML chunks without timing information are treated as high flexible chunks that can be placed "anywhere" in the schedule as long as the assumptions described later are met.
8. The FML chunks in an input document refer to multiple participant and so the scheduler takes into account the possibility of scheduling two chunks of different participants (later transformed into BML blocks) at the same time. The internal representation language allows the scheduler to express this while preparing the document for the transformer.

Internal Representation Language

This language is a hybrid between FML and BML and it is the result of several changes made by the scheduler to the input FML document prior to forwarding it to the transformer.

The FML document declaration section is left unchanged by the scheduler, however the body section, after the scheduling algorithm has finished running, does not contain the three tracks (interactional, performative and mental state) anymore, instead it consists of a single `<body>` element containing several hybrid BML block elements (i.e. `<bml>` elements) in meaningful order.

The BML block elements are hybrid since they contain attributes of the BML language specification (e.g. *id*, *composition*), however they do not embed BML behavior tags but FML function tags instead. We can see this as an initial step that transforms FML chunks (the basic operational unit of our FML specification) into BML blocks (the corresponding unit in BML) while keeping certain functional level information intact (e.g. the track of provenience in the FML document).

The order of appearance of the resulting BML blocks in this hybrid document is meaningful since the scheduler has already sorted out whether the original FML chunks can be scheduled correctly.

The scheduling algorithm described in the following section operates on the FML document with a series of steps and adopts several new attributes and elements that are external to both the FML and BML specifications, thus belonging to our internal representation language.

These attributes are added to the `<fml-chunk>` elements:

- *track:* The track in the document from which the FML chunk comes from;
- *suggestedBMLID:* A suggested value for the *id* attribute of the BML block that will replace the FML chunk later;
- *suggestedComposition:* A suggested value to apply to the *composition* attribute of the BML block created later.

The scheduling of multiple FML chunks at the same time (e.g. with *"start_together"* primitives) is represented with a new tag named `<together>`. This tag has an attribute named *schedule* that can assume two values (*immediate* vs. *normal*). The `<together>` tag is the practical representation of the *Start Together Construct* in the internal representation language, the two values for the schedule *attribute* indicate the construct entry type. The *schedule* attribute is also added to individual FML chunks not belonging to any group and has the same meaning. Thus, this attribute represents the technical solution anticipated earlier to represent the entry type of an arc into a node in the DAG.

Scheduling Algorithm: Implementation

In this section we describe the implementation of the scheduling algorithm. First the `Node` object used to represents nodes in the DAG built by the algorithm is presented. Then the `Constructs Manager`, an helper sub-component in the `FML Scheduler`, is briefly described. Finally, the steps of the algorithm are listed and selected important procedures are detailed in pseudo-code.

The Node Object A `Node` object represents a job in the DAG that is built by our algorithm to decide the schedule. Practically a node corresponds to an FML chunk in the input document. Table 6.4 shows the information included in a Node. There are also FML related information that are not shown in the table but included in the object (*participantRef*, *actID*, *primitive* and *actRef* if any `<timing>` information is provided, *track* type, *suggested BML ID* and *FML functions* embedded in the chunk).

Field	**Type**	**Description**
ID	`Integer`	Unique ID for the node
Visited	`NodeVisitedStatus`	Used by the DFS to establish the status of the node during the visit [*No* | *Temporary* | *Permanently*]
Meeting Time	`Integer`	Meeting time of the node after running the DFS
Farewell Time	`Integer`	Farewell time of the node after running the DFS
ConstructID	`Integer`	Indicates if the node belongs to a *Start Together Construct*, *-1* otherwise
Type	`NodeType`	The type of node [*Root* | *FML_Chunk* | *Start* | *End*]
Entering Type	`NodeEnteringType`	The type of the edge entering the node [*Normal* | *Immediate*]
Parents	`List`	A linked list of pointers to parent nodes
Descendants	`List`	A linked list of pointers to descendants nodes (adjacency list used by DFS)

Table 6.4: The information included in a Node object. These fileds are used by the scheduler algorithm to build the DAG and run the DFS visit on it. The Constructs Manager sub-component also uses some of the information such as the ConstructID and Type. Information related to the FML chunk that the node represents are also incuded but not shown in the table.

The Constructs Manager The *Start Together Construct* can be seen, at an abstract level, as a single special node in the DAG that exposes the embedded *Start* node as entry point for the parent node and the *End* node pointing to descendants. A dedicated sub-component of the `FML Scheduler`, named `Constructs Manager`, aids the creation of those constructs and supports adding a single node to an existing construct by properly updating the node object's information (e.g. the construct ID the node belongs to and the adjacency lists of parent and descendants).

The next page shows the pseudo-code of the *Scheduling Algorithm* divided into 8 steps. The algorithm uses the `Node` object to represent FML chunks (or jobs) in a *DAG*. The `Construct Manager` aids the creation of such DAG abstracting the concept of construct as single "node" to place in the graph and handling all required operations to hold this abstraction (e.g. adding a node to a construct, merging two constructs, etc...) . It also helps holding the assumptions for solving **Problem 2**. The procedures at lines **6.1**, **8.8** and **8.13** are detailed later in pseudo-code.

```
Algorithm 6.5.0: SCHEDULINGALGORITHM
Input Document (FML)
Output: Document (Internal Representation Language)
Data:
  DAG: The directed acyclic graph (initially empty)
  Schedule: List of nodes in the DAG sorted according to timing constraints (initially empty)

Step 1 - Parse input document:
Create an isolated Node for each chunk in the document
Copy chunk information to the node information (track, participant, actID, embedded FML functions)
Create the ROOT node of the DAG

Step 2 - Scan for FML chunks with "immediately" timing:
Create a special Start Together Construct
Make the special Start Together Construct child of the ROOT node of the DAG
for each chunk with "immediately" timing do:
    Add the corresponding node to the special Start Together Construct
    if another node of the same participant is in the construct then assign the node the same suggestedBMLID
    else assign the node a new suggestedBMLID

Step 3 - Scan for FML chunks with "start_together" timing:
for each chunk A with "start_together" timing do:
    if node B specified in A's actRef timing attribute is not in a construct then:
      Create a normal Start Together Construct
      Add A and B to the construct
    else Add the A to the existing construct where B belongs to

Step 4 - Scan for all other FML chunks:
(except those without <timing> information)
for each chunk A do:
    if A has "immediately_after" primitive then:
      Add to DAG following solution to Problem 2 using the Constructs Manager
    else (all other primitives): Add A to DAG as normal node according to primitive
    Update A's entering type

(continue to step 5)
```

(continued)

Step 5 - Handle Isolated Nodes and Constructs:

```
for each isolated Node N or Construct C in the DAG do:
    Make N or C a child of the ROOT node (holding Assumption 2 to solve Problem 2)
    Comment: FML chunks without timing information are not included in the DAG at this stage
```

Step 6 - Obtain Topological Sort of the DAG:

```
Schedule ← TopologicalSort(DAG)
```

Step 7 - Handle Chunks without <timing> Information:

```
for each chunk c without <timing> information do:
    Add the correspondent Node n to the tail of Schedule
```

Step 8 - Transform Input FML Document into Output Internal Document:

```
Create empty output Internal Document
Mirror <declarations> section of input document into output document
Create empty <body> section in output document
Perform two sub-steps working on <body>:
  a) Read nodes from Schedule and write correspondent chunks into Internal Document:
    Comment: chunks corresponding to nodes of the same Start Together Construct
                 are embedded within a pair of <together> </together> tags
    Internal Document ← SchedulerStep8-WriteChunks(Schedule)
  b) Transform FML chunks into BML blocks in Internal Document:
    Comment: <fml-chunk> tags are transformed into <bml> tags having:
                 Attributes: BML 1.0 specification only
                 Content: FML function tags (copied from the original chunks)
    Internal Document ← SchedulerStep8-TransformBlocks(Internal Document)
```

The following is the algorithm's pseudo-code to obtain the topological sorting of the DAG whenever possible, and returning with an error when the topological sort cannot be obtained. The algorithm loops through each node of the graph, starting from the *ROOT* node, initiating a depth-first search that terminates when all nodes have been visited or a node that the DFS is currently visiting (i.e. the DFS is visiting the descendants) is encountered again.

Note that each node n is added to the *VisitedNodesList* only after considering all other nodes in its adjacency list (i.e. all descendants of n in the graph) from left to right, thus holding the assumptions for solving **Problem 2**. Specifically, when the algorithm adds node n to the list, we are guaranteed that all descendants are already in the *VisitedNodesList*: they were added either by the preceding recursive call to *Visit()*, or by an earlier call to *Visit()*.

If no errors are found (i.e. the graph is a DAG) the schedule that the graph represents is **feasible**. Then the algorithm outputs this *Schedule* by ordering the *VisitedNodesList* by farewell time in reverse order (from biggest to smallest) of the nodes included.

```
Algorithm 6.5.1: TopologicalSort(DAG)

main
  global VisitedNodesList ←Empty list that will contain the visited nodes
  comment: Run a DFS of the DAG starting from the ROOT node
  while (there are unmarked nodes)
     do { Select the next unmarked Node n
          (i.e. n's visit status is not marked as "Permanently")
          Visit(n)
  Schedule ← SortByReverseFarewellTime(VisitedNodesList)
  output (Schedule)

procedure Visit(n)
  if n has a "Temporary" mark
    then Stop (Error not a DAG)
  if n is not markerd (i.e. "No" mark)
    then { Mark n as "Temporary"
           Set n's Meeting Time
           for each Node child with an edge from n to child
             do Visit(child)
           Set n's Farewell Time
           Mark n as "Permanently"
           Add n to VisitedNodesList
```

The following is the pseudo-code of the *SchedulerStep8-WriteChunks* procedure. This procedure scans the nodes in the *Schedule* and writes the correspondent FML chunk of each node into the **`<body>`** section of the output *Internal Document*. In this document the order of appearance of **`<fml-chunk>`** elements matters. It is important to remark that FML chunks corresponding to nodes in a *Start Together Construct* (either with *normal* or *immediate* entry type) are embedded within a pair of **`<together>`** and **`</together>`** elements that represent such constructs. This information will be used by the `BML Dispatcher` service later. The **`<together>`** tag also has a *schedule* attribute that can assume either *normal* or *immediate* value depending on the entry type of the construct that it represents.

```
Algorithm 6.5.2: SchedulerStep8-WriteChunks(Schedule)

global InternalDocument
comment: Nodes in Start Together Construct are surrounded by <together> element
comment: Schedule attribute of <together> has entry type of the construct as value
for each (Node n in Schedule)
  do
    Copy corrensponding FML chunk c to <body> of Internal Document
    Add track attribute to c from node n information
    if n is in a Start Together Construct
      then
        Add compositon attribute to c with value "MERGE"
        if n has suggestedBMLID information
          then Add suggestedBMLID attribute to c with value already in n
          else
            if another node of the same participant is in the construct
              then Assign the same suggested BML ID to c
              else Assign the node a new suggestedBMLID
      else
        comment: Chunk c is individual not in a construct
        Add compositon attribute to c with value "APPEND"
        Add suggestedBMLID attribute to c with value already in n
         or suggest a new one
        Add schedule attribute to c according to n's entry type
output (InternalDocument)
```

The following is the pseudo-code of the *SchedulerStep8-TransformBlocks* procedure. This procedure scans the **<fml-chunk>** elements in the **<body>** section of the *Internal Document* and transforms them into **<bml>** elements. These **<bml>** elements have BML 1.0 attributes but they still embed FML functions (copied from the original chunks). The resulting *Internal Document* is ready to be transformed into "pure" BML by the `FML-to-BML Transformer` component according to the XSLT rules.

The *schedule* attribute of **<together>** elements is removed and used to set a *start* attribute that is added to individual functions contained in the **<bml>** elements created. This *start* attribute can be seen as a suggestion for the BML *start* attribute of the behaviors that will replace the functions. FML chunks belonging to a `Start together Construct` with *"normal"* schedule obtain the *start* attribute with value set to "0", whereas when the schedule is *"immediate"* the value is set to "<previous BML BLOCK ID:last BehaviorID in Block:end>" (see marker 1 in algorithm 6.5.3), which practically suggests that the behavior(s) that will replace the function must start at the end of the last behavior of the previous BML block (i.e. immediately after).

In the particular case of individual FML chunks transformed into BML blocks, the *start* attribute is not set (see statements starting at marker 2 in Algorithm 6.5.3). Instead the **<synchronize>** element of BML 1.0 is used[12]. The idea is to synchronize all the functions (later transformed into BML) within such individual chunk to start *"immediately_after"* the end of the last behavior in the previous BML block. This is accomplished by synchronizing all functions in the chunk (marker 4) and starting them after "<previous BML BLOCK ID:last BehaviorID in Block:end>" (marker 3).

FML chunks having the same *suggestedBMLID* value are merged into single **<bml>** elements, i.e. conceptually the functions in the chunks with the same *suggestedBMLID* are placed in the same BML block.

---

[12] See: `http://www.mindmakers.org/projects/bml-1-0/wiki#ltsynchronizegt`

```
Algorithm 6.5.3: SchedulerStep8-TransformBlocks(InternalDocument)

global InternalDocument
for each FML chunk c in <body> of InternalDocument
  do
    if c has suggestedBMLID not seen before
      then
        Create <bml> element b
        Add characterID attribute to b with value participantRef of c
        Add id attribute to b with value suggestedBMLBlockID of c
        Add composition attribute with value suggestedComposition of c
        Copy all functions of c to the newly created element b
      else
        comment: Merge functions into one BML Block
        Append all functions of c in existing BML element b
        with same id as suggestedBMLID of c
    for each function f in chunk c and copied to b
      do
        comment: Copy track value from the chunk
        Add track attribute of c to f
        comment: Add the start value to each function
        if c belongs to a Start Together Construct
          then
            if construct schedule value is "normal"
              then Set start = "0"
            else if construct schedule value is "immediate"
              then
                comment: Start immediately after                                  (1)
                Set start = "<previous BML block ID:last BehaviorID in block:end>"
          else
            comment: Chunk c not in a construct
            if c schedule value is "normal"Set start = "0"
            else if c schedule value is "immediate"
              then
                comment: Synchronize all functions
                start is not set                                                  (2)
                Add a <constraint> element to the BML block b
                Add a <after> element as child of <constraint>
                Set ref attribute of <after> to
                "<previous BML BLOCK ID:last BehaviorID in Block:end>"            (3)
                Add <sync> element as child of <after> for each function f
                Set ref attribute of <sync> to "<function id:start>"              (4)
output (InternalDocument)
```

Further Remarks

When the scheduling algorithm encloses FML chunks within the same `<together>` element (originally `Start Together Construct`) and these chunks obtain the same *suggestedBMLID* value (e.g. when they refer to the same *participant*), the algorithm follows the design principle consisting of putting as much as possible into the same BML block. In fact chunks with the same BML ID are merged into the same `<bml>` element in later stages of the algorithm. In **step 2** they receive the *suggestedBMLID* value and in **sub-step b** of **step 8** chunks having the same *suggestedBMLID* are merged.

In **step 4**, when encountering an FML chunk A having as primitive A *must_end_before* B. There are three scenarios that the algorithm handles:

- B has *"immediately"* timing primitive. In this case the algorithm returns an error since the schedule is not feasible.
- B has to *"start_immediately_after"* C. In this case A is scheduled to *"start_together"* with C (parent node in the DAG and referred by B's primitive).
- In all other cases B becomes a child of A in the DAG.

Every operation on the DAG built by the algorithm ensures that the `Start Together Constructs` are kept at the node level of abstraction. This means that all operations involving a construct deal with the *Start* and *End* nodes of the construct without having to deal with the inner nodes.

The `Construct Manager` takes care of these operations and makes sure to hold the assumptions for solving **Problem 2**. As an example, consider the case where a node A represents an FML chunk that is scheduled to *"start_sometime_after"* B, where B is in a construct. The `Construct Manager` in this case ensures that an edge between A and B in the graph is established by linking the node A (child) in the adjacency list of the *End* node in the construct where B belongs to. Operations involving the inclusion of a node in an existing construct or merging two existing constructs are also supported by the `Construct Manager`.

Scheduling Examples for Two Scenarios

In this section we continue with the two full FML examples described in the previous chapter, in Section 5.3.4, and we show how these FML documents are handled by the `FML Scheduler` when received as input.

Since the declaration section of an input FML document is not used at this stage, we will omit it in the examples shown. We first show the *Internal Document* that is output by the *SchedulerStep8-WriteChunks* procedure during **step 8** of the scheduling algorithm, when tracks are removed and the **`<fml-chunk>`** elements appear in meaningful order. Then we show the *Internal Document* that the scheduling algorithm outputs prior to becoming input to the `FML-to-BML Transformer` component.

Ordering a Cheeseburger According to the temporal constraints specified in the FML document referring to the cheeseburger example at page 122 of Section 5.3.4, the FML chunks should be scheduled as follows:

ACT01 → ACT02 ⇸ ACT04 ⇸ (ACT03 ↔ ACT05) ⇸ ACT06

We can see that ACT04 is scheduled *"immediately_after"* ACT02. Whereas ACT03 and ACT05 are scheduled together in a `Start Together Construct` that, in turn, is scheduled to start *"immediately_after"* ACT04. Finally, ACT06 is scheduled *"immediately_after"* the construct.

The FML chunks corresponding the schedule shown above are listed as follows in the *Internal Document* adopting the internal representation language.

```
1  <body>
2
3    <fml-chunk actID="ACT01" participantRef="idPete" schedule="normal" track="interactional"
4              suggestedBMLBlockID="BML-001" suggestedComposition="APPEND">
5          <grounding floorID="floor1" id="id1" type="ack" />
6      </fml-chunk>
7
8      <fml-chunk actID="ACT02" participantRef="idPete" schedule="normal" track="interactional"
9                suggestedBMLBlockID="BML-002" suggestedComposition="APPEND">
10         <turn-taking floorID="floor2" id="id2" type="take" />
11         <speech-act floorID="floor2" id="id3" type="request" />
12     </fml-chunk>
13
14     <fml-chunk actID="ACT04" participantRef="idPete" schedule="immediate" track="performative"
15               suggestedBMLBlockID="BML-003" suggestedComposition="APPEND">
16         <performative-extension id="id5" floorID="floor2" addressee="idGeorge" >
17             <discourse-structure type="topic">
18             George make a <rhetorical-structure type="emphasis">cheesburger</rhetorical-structure>
19         </discourse-structure>
20         </performative-extension>
21     </fml-chunk>
22
23     <together schedule="immediate">
24         <fml-chunk actID="ACT03" participantRef="idPete" schedule="immediate" track="interactional"
25                   suggestedBMLBlockID="BML-004" suggestedComposition="MERGE">
26               <turn-taking floorID="floor2" id="id4" type="give" />
27           </fml-chunk>
28
29           <fml-chunk actID="ACT05" participantRef="idPete" schedule="immediate" track="mental-state"
30                     suggestedBMLBlockID="BML-005" suggestedComposition="MERGE">
31           <emotion floorID="floor2" id="id6" type="anger" regulation="fake" intensity="0.7"
32                   weightFactor="1.0"  />
33           </fml-chunk>
34     </together>
35
36     <fml-chunk actID="ACT06" participantRef="idPete" schedule="immediate" track="mental-state"
37               suggestedBMLBlockID="BML-006" suggestedComposition="APPEND">
38         <emotion floorID="floor2" id="id7" type="anger" regulation="fake" weightFactor="0.0" />
39     </fml-chunk>
40
41 </body>
```

FML Listing 6.1: The *Internal Document*'s body section of the cheeseburger example after that the *SchedulerStep8-WriteChunks* procedure of the scheduling algorithm is executed.

In the *Internal Document* shown above, the FML chunks contrary to the original *FML Document* appear in a meaningful order. The value of the *schedule* attribute indicates whether each chunk needs to be scheduled immediately after the previous (*immediate*) or not (*normal*). The *track* has become an attribute of the **`<fml-chunk>`** element.

The chunks identified with *id* values ACT03 and ACT05 appear in a **`<together>`** element representing a `Start Together Construct` and indicating that the two chunks are scheduled to start together (see line 23). The **`<together>`** element also has the *schedule* attribute indicating that it has to be scheduled immediately after the previous chunk (ACT04).

Every FML chunk also has two additional attributes, *suggestedBMLBlockID* and *suggestedComposition*, that will be used by the second and final sub-step of the scheduling algorithm (step 8) as we will see in the example that follows. The functions inside each chunk are copied from the original *FML Document* as they appear.

The next example shows the *Internal Document* that the `FML Scheduler` outputs after the second sub-step in **step 8**.

```
<body>

  <bml characterId="idPete" id="BML-001" composition="APPEND">
    <grounding floorID="floor1" id="id1" type="ack" track="interactional" start="0" />
  </bml>

  <bml characterId="idPete" id="BML-002" composition="APPEND">
      <turn-taking floorID="floor2" id="id2" type="take" track="interactional" start="0" />
      <speech-act floorID="floor2" id="id3" type="request" track="interactional" start="0" />
    </bml>

  <bml characterId="idPete" id="BML-003" composition="APPEND">
    <performative-extension id="id5" floorID="floor2" addressee="idGeorge" track="performative">
      <discourse-structure type="topic">
          George make a <rhetorical-structure type="emphasis">cheesburger</rhetorical-structure>
      </discourse-structure>
    </performative-extension>

    <constraint>
      <after ref="BML-002:id3:end">
        <sync ref="id5:start" />
      </after>
    </constraint>
  </bml>

  <together>
    <bml characterId="idPete" id="BML-004" composition="MERGE">
      <turn-taking floorID="floor2" id="id4" type="give" track="interactional" start="BML-003:id5:end" />
    </bml>

    <bml characterId="idPete" id="BML-005" composition="MERGE">
      <emotion floorID="floor2" id="id6" type="anger" regulation="fake" intensity="0.7" weightFactor="1.0"
              track="mental-state" start="BML-003:id5:end" />
    </bml>
  </together>

  <bml characterId="idPete" id="BML-006" composition="APPEND">
    <emotion floorID="floor2" id="id7" type="anger" regulation="fake" weightFactor="0.0" track="mental-state" />

    <constraint>
      <after ref="BML-005:id6:end">
        <sync ref="id7:start" />
      </after>
    </constraint>
  </bml>

</body>
```

FML Listing 6.2: The body section of the cheeseburger example produced as the output at the end of the scheduling algorithm.

In the final *Internal Document* output by the algorithm and shown above, the **`<fml-chunk>`** elements are replaced by **`<bml>`** elements. The *participant* attribute is replaced by the *characterId* one, the values for the *id* and *composition* attributes are taken, respectively, from the *suggestedBMLID* and *suggestedComposition* attributes of the FML chunks.

The functions inside the BML blocks are still represented in FML but they have the *track* attribute taken from the chunk of origin. This value can be used later during the transformation process.

The *schedule* attribute does not appear in the document anymore, instead the *start* attribute is added to each function according to the values of the *schedule* attributes in the original chunks or `<together>` elements.

We can see, for example, how these values are set for the functions at line 8. Furthermore, the last BML block with *id* set to "BML-006" has a `<constraint>` element inside (see line 19). This BML 1.0 element says that the behavior with *id* set to "id7" (now a function but later will be transformed into a behavior) needs to start after the behavior with *id* set to "id6" in the block "BML-005" ends (i.e. immediately after the end of the last behavior of the previous BML block).

The Very Beginning of a Virtual Greeting Encounter This second example shown in the original *FML Document* at page 127 of Section 5.3.4 specifies the following schedule:

$$\text{ACT01} \nrightarrow \text{ACT02} \rightarrow \text{ACT03}$$

We can see that ACT02 is scheduled *"immediately_after"* ACT01 and ACT03 after ACT02 with *normal* entry type.

The *Internal Document* produced by the *SchedulerStep8-WriteChunks* procedure and including the scheduled FML chunks of the greeting example is shown in the FML listing 6.3:

```
<body>

  <fml-chunk actID="ACT01" participantRef="idAgent" schedule="immediate" track="interactional"
            suggestedBMLBlockID="BML-001" suggestedComposition="MERGE">
      <initiation floorID="floor1" id="id1" type="recognize" addressee="idUser" />
    </fml-chunk>

    <fml-chunk actID="ACT02" participantRef="idAgent" schedule="immediate" track="interactional"
                suggestedBMLBlockID="BML-002" suggestedComposition="APPEND">
        <initiation floorID="floor1" id="id2" type="salute-distant" addressee="idUser" />
      </fml-chunk>

    <fml-chunk actID="ACT03" participantRef="idAgent" schedule="normal" track="mental-state"
              suggestedBMLBlockID="BML-003" suggestedComposition="APPEND">
        <cognitive-process floorID="floor1" id="id3" type="idle" weightFactor="0.8" />
  </fml-chunk>

</body>
```

FML Listing 6.3: The *Internal Document*'s body section of the greeting example after the *SchedulerStep8-WriteChunks* procedure of the scheduling algorithm is executed.

Finally, the *Internal Document* that the scheduling algorithm produces as output and sends to the `FML-to-BML Transformer` is shown in listing 6.4.

```
<body>

  <bml characterId="idAgent" id="BML-001" composition="MERGE">
      <initiation floorID="floor1" id="id1" type="recognize" addressee="idUser" track="interactional"
                start="0" />
  </bml>

  <bml characterId="idAgent" id="BML-002" composition="APPEND">
      <initiation floorID="floor1" id="id2" type="salute-distant" addressee="idUser" track="interactional" />

      <constraint>
        <after ref="BML-001:id1:end">
          <sync ref="id2:start" />
        </after>
      </constraint>
  </bml>

  <bml characterId="idAgent" id="BML-003" composition="APPEND">
    <cognitive-process floorID="floor1" id="id3" type="idle" weightFactor="0.8" track="mental-state"
      start="0" />
  </bml>

</body>
```

FML Listing 6.4: The body section of the greeting example produced as output at the end of the scheduling algorithm.

6.5.4 FML to BML Transformer

This component receives transformation requests from the `FML Scheduler` of *Internal Documents* written using the internal representation language (see Figures 6.2 and 6.5 for a visual reference of the service module's architecture).

The transformer initially loads an XSLT style sheet that allows a designer to specify the transformation rules. These rules are meant to be applied to FML function tags within each BML block included in the *Internal Document* received as input. The full set of rules is available in Appendix E.

In our implementation we focused on the specific context of greeting encounters, therefore the rules that we specified operate exclusively on functions in the *Interactional* track, in particular the *Initiation* and *Closing* categories. For the *Initiation* category our rules handle the following functions: *react*, *recognize*, *salute-distant*, *approach-react*, *salute-close*, *initiate*. Whereas for the *Closing* category we created rules for the *break-away* and *farewell* functions[13].

The rules make use of the contextual information provided by the *Internal Document*'s declaration section, and depending on the participant's attitude (affiliation) and personality (extraversion) they establish which multimodal behaviors can accomplish the functions in the document. The parameters of these behaviors are also modulated (e.g. gaze behavior duration). These rules basically follow the theoretical framework described in Chapter 4.

The transformer, after the application of the rules, outputs a transformed document containing BML blocks (i.e. **`<bml>`** elements) that might refer to multiple participants (characters in the BML terminology). All the tags and attributes of this document are described in the BML 1.0 specification except for the **`<together>`** elements introduced by the `FML Scheduler` and kept unchanged by the transformer. These elements support the realization of BML blocks belonging to multiple characters and will be processed later by the `BML Dispatcher` component.

---

[13] All other FML functions introduced with our specification and without any matching transformation rule are handled by a general XSLT rule that converts the functions into BML **`<wait>`** elements (i.e. no operation).

Example: Very Beginning of a Virtual Greeting Encounter

This example shows the output of the `FML-to-BML Transformer` for the *Internal Document* that is received by the scheduler representing the virtual greeting encounter scenario. The transformer dumps into the output document in the form of comments the declaration section of the input document and the original FML functions prior transforming too.

```
1  <saiba-act>
2  <!--
3  DECLARATIONS
4
5  IDENTIKITS:
6  1) id [idAgent] name [Agent] gender [male] personality(extraversion) [HIGH]
7    relationship [ACQUAINTANCE with idUser]
8  2) id [idUser] name [User] gender [female]
9
10 FLOORS:
11 1) floorID [floor1] floor-cfg [unicast]
12   PARTICIPANTS:
13   1) identikitRef [idAgent] role [speaker] entity [individual] attitude(affiliation) [FRIENDLY towards idUser]
14   2) identikitRef [idUser] role [addressed-hearer] entity [individual]
15 -->
16
17 <bml characterId="idAgent" id="BML-001" composition="MERGE" xmlns="http://www.bml-initiative.org/bml/bml-1.0">
18
19   <!-- tag [initiation] floorID [floor1] ID [id1] type [recognize] addressee [idUser] track [interactional] -->
20
21   <gaze id="id1_gaze_1" start="0" end="id1_gaze_1:start + 3" influence="HEAD" target="idUser" />
22   <head id="id1_headtoss" start="id1_gaze_1:start + 0.6" end="id1_headtoss:start + 0.8" lexeme="TOSS" />
23   <faceLexeme id="id1_raisebrows" start="id1_headtoss:start" end="id1_headtoss:end" a
24         attackPeak="id1:start + 0.4" lexeme="RAISE_BROWS" amount="0.5"/>
25 </bml>
26
27 <bml characterId="idAgent" id="BML-002" composition="APPEND" xmlns="http://www.bml-initiative.org/bml/bml-1.0">
28
29   <!-- tag [initiation] floorID [floor1] ID [id2] type [salute-distant] addressee [idUser]
30       track [interactional] -->
31
32   <postureShift id="id2_postureShift" start="0">
33     <stance type="STANDING"/>
34     <pose part="WHOLEBODY" lexeme="FACE" fga:target="idUser" xmlns:fga="http://cadia.ru.is/FMLGreetingAgent"/>
35   </postureShift>
36
37   <gazeShift id="id2_gaze_1" start="0" influence="HEAD" target="idUser"/>
38   <gazeShift id="id2_gaze_2" start="id2_gaze_1:start + 2" influence="HEAD" target="idUser"
39         offsetAngle="45.0" offsetDirection="DOWNLEFT"/>
40
41   <gesture id="id2_gesture" start="id2_gaze_1:start + 1" lexeme="SHORT-WAVE" mode="RIGHT-HAND"/>
42
43   <faceLexeme id="id2_face_1" start="id2_gaze_1:start" end="id2_face_1:start + 30"
44         attackPeak="id2:start + 8" overshoot="20" lexeme="RAISE_BROWS" amount="0.8"/>
45   <faceLexeme id="id2_face_2" start="id2_gaze_1:start" end="id2_face_2:start + 30"
46         attackPeak="id2:start + 8" overshoot="20" lexeme="RAISE_MOUTH_CORNERS" amount="0.8"/>
47
48   <constraint>
49     <after ref="BML-001:id1:end">
50       <sync ref="id2_postureShift:start"/>
51       <sync ref="id2_gaze_1:start"/>
52     </after>
53   </constraint>
54 </bml>
55
56 <bml characterId="idAgent" id="BML-003" composition="APPEND" xmlns="http://www.bml-initiative.org/bml/bml-1.0">
57   <!-- tag [cognitive-process] floorID [floor1] ID [id3] type [idle] track [mental-state] weightFactor [0.8] -->
58   <wait id="id3" start="0" duration="0.1"/>
59 </bml>
60
61 </saiba-act>
```

FML Listing 6.5: The BML generated as output by the `FML-to-BML Transformer` for the virtual greeting encounter example.

The example above shows the BML produced by the `FML-to-BML Transformer` for the “recognize” and “salute-distant” functions accomplished by an agent during a greeting encounter with the user. The “recognize” function is accomplished with a gaze behavior towards the user (line 21) and an additional head toss (lines

22-24) due to the relationship level between the twos (set to "ACQUAINTANCE", see line 7).

The agent manages a friendly impression of affiliation by producing a smiling facial expression (lines 43-46), in addition to a gaze towards the user (line 37) and a short hand wave (line 41) when accomplishing the "distant salutation" function. A posture shift (lines 32-35) allows the agent to orient itself towards the user. The **`<pose>`** element inside the **`<postureShift>`** has a newly defined attribute with the corresponding namespace (*fga*) as suggested in these cases by the BML 1.0 specification.

6.5.5 BML Dispatcher

The `BML Dispatcher` is the last component in the pipeline. It carries out a simple task, it reads an incoming BML document containing BML blocks (i.e. **<bml>** elements) referring to multiple agents and dispatches each block to the correct agent module as indicated by the *characterId* attribute (originally the participant in FML).

Dispatching a BML block simply means sending it to the agent module in order to be executed by the `BML Realizer`. A basic multi-character synchronization is obtained with **<together>** elements. The dispatcher strips out these element from the BML document when it encounters them and sends in parallel the embedded BML blocks to the respective agents (by using a separate thread for each BML block).

The syntax of the BML documents generated by the service module and processed by the `BML Dispatcher` component, except for **<together>** elements, has been validated against the BML 1.0 specification[14].

---

[14] We used the XSD schema available here: `https://github.com/saiba/BMLxsd`

6.6 Application in the Virtual Reykjavik Learning Environment

The agent and service modules have been deployed in the "Icelandic Language and Culture Learning in Virtual Reykjavik" project, the 3D learning environment introduced in Section 1.2. Currently a 3D reconstruction of a square named *Austurvöllur* and located downtown Reykjavik is featured in the project.

Figure 6.10 shows a screenshot of the square with an `FML Greeting Agent` performing the distant salutation while greeting the user approaching in first person perspective.

Figure 6.10: A screenshot taken from the "Icelandic Language and Culture Learning in Virtual Reykjavik" project. The female `FML Greeting Agent` featured in the scene is performing the distant salutation towards the user that is approaching her in first person perspective view.

The project uses the Unity3D game engine[15], thus the agent and service module have been developed as two C# scripts. Every relational agent in the square provided with impression management capabilities has the `FML Greeting Agent` module's script attached. The `FML Service` script runs centrally.

---

[15] See: http://unity3d.com

Both scripts can be plugged right away into other ECA systems that adopt the Unity3D game engine. While the service script has been kept as much as possible away from unity related features in order to be easily re-implemented in a different programming language for another game engine, the agent module is more Unity dependent.

In particular, the service module only requires XML related features that are available in all common programming languages and frameworks (XML documents parsing, validation and XSLT-based transformations). On the other hand, the `Perception` and `Input Understanding` components of the agent module use a continuous *update* loop provided by Unity scripting mechanics (in combination with vectors) and *Spehere Colliders*[16] respectively, to detect the raw user-agent distance and to model the *React Area* described earlier.

The virtual characters in the screenshot have been modeled with the Autodesk's *Pinocchio* project web tool[17].

---

[16] See: `docs.unity3d.com/Documentation/Components/class-SphereCollider.html`

[17] See: `projectpinocchio.autodesk.com`

6.7 Evaluation

6.7.1 Model Evaluation

The implementation of the two modules presented in this chapter thoroughly followed the criteria defined in our theoretical framework. The final application in Virtual Reykjavik represents a particular instance of our framework that has been already evaluated by users with the three studies described earlier.

The models and the empirical evidence obtained have been reflected in every stage of the modules' design and implementation. The `Perception` and `Input Understanding` components of the agent module detected approaching users and decoded functional meanings of those raw distances in the same fashion as was done for all the three evaluation studies conducted earlier. The greeting communicative functions planned by the `Reactive Planner` have been mapped to multimodal behavior selected on the basis of our discoveries. These behaviors and the parameters that made them distinct were generated to manage impressions of personality and attitude on the users.

This practical application represents a common place where all the discoveries flow into. However, we are aware that some distortions of the original model could arise when moving from theory to practice, in particular considering that in the final implementation some of the constraints imposed in the previews experimental setting have been relaxed. For example, while in two of the evaluation studies users merely observed agents' pre-scripted reactions, in the Virtual Reykjavik application users needed a higher level of control with more freedom of movement while approaching the agents. Furthermore, the agents planned and accomplished greeting communicative functions in real time. Implications of this aspect will be discussed in the last section of this chapter along with some limitations of our solution.

6.7.2 Technical Evaluation

By making our modules SAIBA compliant we sought to obtain flexibility of integration for our solution into existing fully working SAIBA systems. In addition to this design goal, we think that scalability of our solution is also an important performance issue considering that the service module is meant to be a shared central solution to allow the agents of a system to manage their impressions on users in real time. These two aspects are discussed in the following sections.

Flexibility

Other SAIBA compliant systems adopting Unity as 3D engine can easily benefit from our solution. The only issue that might slow the integration process might be represented by the accessibility to a BML Realizer. When available, our modules need only to be configured with the public parameters shown earlier, whereas when not available, programming and animation work is needed to execute the BML output from our service module.

The migration of our solution to another 3D engine is more complex but also possible. It requires the (re)development of the two modules in the programming language of the target engine, furthermore it needs the adaptation to the new engine of the agent module parts that rely on Unity3D features described earlier.

Performance

The `Perception` component in the agent module relies on Unity3D native features that have already been subject to testing and the behavior realization (either via standard procedural animation techniques or external BML Realizers) does not represent a crucial part of our solution to test. In fact, the agent module supports the presence of a third party BML Realizer that we can consider external to our solution when measuring performance.

The service module, instead, is mostly game engine independent and reflects some design choices of our own that require a performance evaluation concerning scalability, considering that it has been designed to be a central shared solution accessible from multiple agents in real time. In particular, its SAIBA compliance involves the usage of XML-like languages, such as FML and BML, that might lead to processing time issues when performing a high number of operations including parsing documents (in our case it might be when the `FML Scheduler` parses an input *FML Document*) or generating documents (e.g. the transformation operated by the `FML-to-BML Transformer`).

These aspects could clearly impact the performance of our solution. Therefore we focused on scalability by conducting a stress test of the service module to validate how it reacts to a growing number of incoming requests in a scenario where multiple (fictitious) `FML Greeting Agents` are registered to the service (i.e. are active in a system) and are sending transformation requests at different rates.

We were not interested in performance related to 3D rendering or behavior realization, therefore we deployed an empty test environment without the necessity to render any geometry (e.g. virtual characters) in order to focus our test exclusively on the `FML Service` performance.

In particular we tested the **average time** required by the module to process all incoming transformation request in a given **time interval**. The processing time for a *single request* was computed by adding up the time taken by the `FML Scheduler` to process an input FML document and the time taken by the `FML-to-BML-Transformer` component to transform the generated *Internal Document* containing the scheduled FML into BML blocks ready to be dispatched. We excluded the `BML Dispatcher` processing time from this computation since it does not accomplish complex operations.

As input FML document for the test we used the greeting example 5.7 shown at page 129 in Section 5.3.4. We had a growing variable number of fictitious `FML Greeting Agents` ($\#AGENTS = 1, 10, 30, 40, 50, 100, 1000$) sending transformation requests of such FML document to the service module. We ran different test cases where, given a number of agent modules, every module was operating at the same frequency, i.e. was sending a transformation request every *REQUEST FREQUENCY* milliseconds ($REQUEST\ FREQ. = 300ms, 200ms, 100ms$). We also had a test case were the frequency was random. This allowed us to to model a more realistic testing scenario where each agent operated at a different random frequency in the interval $[0, 300]$ milliseconds. Each test case was run for a **time interval** of **60 seconds** and afterwards the average processing time of all the requests sent was computed.

We used a desktop computer equipped with an AMD Athlon II X2 260 Processor operating at 3.20 GHz, with 2.00 GB of RAM and running Windows 7 64 bit as operating system. Table 6.5 shows the results of these tests. For each test case identified with the number of agents (see the *NUMBER of FML GREETING AGENT MODULES* column) operating at a given request frequency (see the column *REQUEST FREQ.*) we show the average time (in seconds) taken by the service module to process the incoming requests.

REQUEST FREQ. (in ms)	**NUMBER of FML GREETING AGENT MODULES**						
	1	10	30	40	50	100	1000
300	0.12	0.15	0.23	0.28	0.34	15.25	31.15
200	0.14	0.15	0.26	0.4	3.59	18.64	34.93
100	0.15	0.17	4.01	9.63	17.38	22.06	27.81
RANDOM [0,300]	0.15	0.16	0.57	4.47	8.58	18.14	31.93

Table 6.5: Summary of the stress test. We had a different number of fictitious `FML Greeting Agents` ($n = 1, \dots, 1000$) sending FML transformation requests to the `FML Service` module. Every agent module was operating at a frequency indicated in the first column, i.e. was sending a transformation request every *REQUEST FREQ.* milliseconds. The values in correnspondence of each *REQUEST FREQ.* and *NUMBER of AGENTS* indicate the average time (in seconds) taken by the service module to serve all the incoming request after 60 seconds of operation.

From these results we can observe how the growing number of agents and the higher frequency (a smaller *REQUEST FREQ.* value indicates a shorter interval between two consecutive requests sent by an agent) lowers the service module's performance. In particular, when 30 agents are sending each one a request every 100ms we have a significant drop in performance with the service module taking an average of 4.01 seconds to process each incoming request.

Either when the agents are operating at fixed or random frequencies (i.e. in a scenario closer to the real application of our modules), we can see that the average response time is within the range 0.12s - 0.17s (120ms - 170ms) with the number of agents ranging from 1 to 10. These are acceptable performance for exhibiting believable behavior in a human-agent interaction scenario compared to the variation in human responses in human-human interactions (in the range 100ms - 300ms) as suggested by [Thórisson, 1994]. While this number of agents (10) might represent an important scalability issue in different application domains (e.g. crowd simulations), it is still a desirable target in our application domain. In fact, on the basis of our theoretical framework the two modules are meant to model scenarios were a *single* greeting relational agent is managing impressions on the user. Further insights on the applicability of our solution to a growing number of agents will be provided in the next section. These aspect not only involve adjustments to the design and implementation of the modules, but also further research on a theoretical basis to model such new scenarios.

6.8 Final Considerations and Limitations

This chapter presented the design principles, the implementation details, an application and the technical evaluation of a novel SAIBA compliant computational solution aimed at providing a relational agent with impression management capabilities during a first greeting encounter with the user (modeling 1-to-1 interactions). Our solution includes two software modules named `FML Greeting Agent` and `FML Service`. The former module needs to be plugged into each agent in a system and provides perception, communicative greeting functions planning and behavior realization capabilities. The latter is a central shared module that provides scheduling, transformation and delivery services of input documents consisting of FML chunks representing communicative functions transformed into output documents of BML multimodal behavior blocks. The two modules are meant to be integrated into fully working systems that already provide interactional and conversational capabilities once the greeting encounter ends and the conversation with the user begins. However, the service module also represents a data-driven domain independent FML processor that could be easily adopted from different SAIBA compliant multi agent systems to schedule and transform the agent's communicative functions into multimodal behavior.

The solution presented puts into practice the notions of our theoretical framework presented in Chapter 4 and it represents the first SAIBA compliant work featuring the FML specification that we proposed in Chapter 5. An application into the upcoming system "Icelandic Language and Culture Training in Virtual Reykjavik" has been presented in Section 6.6.

6.8.1 Limitations

There are some limitations of the proposed solution that we discuss in this section along with possible ways of solving them.

Model and Design Issues

The theoretical framework that we introduced is based on a model of 1-to-1 human greeting interactions and so our evaluation studies were designed to encompass single user-agent interactions. These theoretical foundations shaped the design of our agent and service modules accordingly allowing an agent to manage impressions on individual approaching users. Furthermore, the framework doesn't deal with the farewell dynamics, therefore we implemented a basic solution that detects when the user moves away to accomplish the break-away and farewell functions. This model related issue will be further discussed in the following chapters together with further insights related to what needs to be done to enhance the implementation of such dynamics (e.g. farewell) in our modules.

Implementation Issues

Error management in the agent and service modules. The current implementation handles major errors by logging them. However, there are some cases where further management might be a convenient feature to implement. For example, the `FML Scheduler` decides whether a given FML document produces a feasible schedule, but doesn't propose a suggestion to fix it in case of scheduling conflicts or inconsistencies.

Perception component in the agent module. The current perception, due to theoretical foundations, has been designed to support the detection of a single *user* approaching the agent that runs the agent module. However, there are two scenarios that the agent module currently doesn't handle. First, when another *user* (e.g. in a multi-player scenario this could be another agent in the system or a real user) enters the *React Area* and there is already the current approaching *user* inside that the agent was reacting to. In such scenario, the current approaching *user* might move towards other directions (still in the area) or might face elsewhere, thus terminating the greeting process earlier and giving the agent the opportunity to greet the other *user* entering the area. Secondly, the design and implementation of mechanisms to negotiate the reopening of a greeting process aborted earlier than the final phase of the greeting (for example the current approaching *user*, in the previously described scenario, moving again towards the agent) are needed. The current implementation only avoids having an agent perform the same communicative function multiple times during a greeting, but "reopening" a greeting involves more sophisticated dynamics that might not only require future work at design and implementation levels, but also on a theoretical level.

FML scheduling in the service module. The current `FML Scheduler` implementation assumes that the input FML document has unchangeable temporal constraints, thus a schedule can be only validated but changes to the timing primitives are not allowed (as discussed earlier). Furthermore, the scheduler doesn't support plan adaptations on-the-fly, i.e. it doesn't allow a new incoming FML document to change a previously submitted one prior to getting transformed into BML in scenarios where the `Reactive Planner` wants to modify the planning of a given set of FML chunks previously made. A similar problem has been encountered with BML scheduling and [Reidsma et al., 2011] proposed a solution involving the concept of "dependent" BML blocks (see the paper for further details).

Transformation rules in the service module. The XSLT rules to transform FML into BML only focus on functions in the interactional track of the "initiation" and "closing" categories. Further rules are needed to map all the functions proposed in our FML specification onto multimodal behavior.

Ease of integration. We chose Unity3D as game engine due to its popularity and since the "Virtual Reykjavik" learning environment was available under that framework. The integration of our modules into systems adopting other graphic engines requires the re-implementation of the modules in the supported scripting languages (e.g. other engines such as Ogre3D[18] or Unreal[19] adopt, respectively, C++ or Unreal Scripting Language).

---

[18] See: `www.ogre3d.org`

[19] See: `www.unrealengine.com`

SAIBA Issues

Multi-character systems and BML. Even though BML provides a clear-cut specification of the internal multimodal synchronization of the behavior of an agent (i.e. character or participant), it lacks the expressiveness to specify the interaction of this behavior with other agents. We provided a simple mechanism to provide multi-character synchronization with the `<together>` element similarly to [Aggarwal and Traum, 2011] where they used an additional BML event. While our solution allows the service module to schedule chunks at the same time (that later become BML blocks), more sophisticated control mechanisms are required. Suggestions for BML extensions that address these shortcomings have been introduced in [Zwiers et al., 2011, Reidsma et al., 2011, van Welbergen et al., 2012] and we also argue that the SAIBA framework, and the BML specification in particular, should be augmented with such control mechanisms.

BML expressiveness. In the greeting example 6.5 we showed the BML generated by the `FML-to-BML Transformer` component. In our evaluation studies, the smiling behavior (among the others) accomplishing the friendly "distant salutation" function was gradually starting when the user was at the "distant salutation" distance and continuing till the "initiation" of the conversation. Since BML does not currently support this continuity (or spanning of a behavior across multiple blocks) we hard-coded this feature in our pseudo-realizer. The BML *composition* attribute has been introduced to handle a new realization request of a BML block that is sent before the realization of previous requests complete. However, none of the possible values for this attribute[20] allowed us to achieve the desired continuous realization for that behavior.

BML realization. The issue mentioned above involves another SAIBA related problem concerning the interchangeability of BML realizers implemented by different research groups. A particular instance of a BML realizer could yield to different execution manners of the same behavior specified with the same BML attributes compared to another instance. This would suggest to the SAIBA community the revision of the BML specification in order to provide a realizer with more information concerning not only the behavior to execute but more technical

[20] Possible values are: *MERGE*, *APPEND* or *REPLACE*.
See here for more information: http://www.mindmakers.org/projects/bml-1-0/wiki/Wiki#Composition

details on selected parameters that allows an agent to realize a behavior in the same manner across different realizer implementations.

Technical Issues

Graphics. There are some graphical issues concerning the visibility of the agents' exhibited nonverbal behavior that depend on the particular environment where the modules are used. For this thesis we applied them to agents in a 3D virtual environment displayed on a regular monitor. The user can engage in greetings with the relational agents in the environment by navigating it in a first person perspective camera view as opposed to the public space scenario described in Section 4.4, where the user approaches the agent in the real world. We identified three main issues that sometimes affected the visibility of some behavior in our selected application. These issues are: (1) the size and resolution of the display, (2) the character's model detail and (3) the virtual distance between the agent and the user (i.e. the camera). For example, it might be hard to observe facial expressions or gaze behavior involving short eye-only movements when the virtual character is quite far from the user or the character's face has a poor level of detail. We faced the distance issue, for example, when our agents were performing smiling and gazing behavior to accomplish the distant salutation function. The smiling behavior was only visible when the user was closer to the agent and after it had already started. We tweaked some of the camera parameters to fix this problem (i.e. the field of view), but other solutions might involve more sophisticated camera techniques, for example zooming at the occurrence of specific agent behavior that must be observed by the user. The usage of larger displays perhaps with larger resolution might be another possible solution.

Performance. The service module's performance drops down when more than 10 agents are sending simultaneous transformation requests at random rates. While we argued that this issue doesn't represent a big concern for the application of our modules, it still remains desirable to perform better in terms of scalability considering that the FML specification and the scheduling and transformation services of the module could be applied to different domains and agent systems in the future. The current service module design supports a separate thread for each component of the service module. Each service has its own incoming requests queue. A possible improvement consists of assigning each incoming requests, for a given component, to a separate thread. Therefore multiple requests (for a given component) could be handled in parallel in this way. However, one has to make sure that requests coming from the same agent module are handled in the order they have been received.

"Understanding human language, imperfect and at the same time capable of realizing a supreme imperfection that we call poetry, represents the only conclusion of every pursuit for perfection."

Umberto Eco (1932 – present)

7

Conclusions

This thesis presented a complete framework for analyzing and modeling human nonverbal communicative behavior for virtual relational agents (and embodied conversational agents in general) to provide them with impression management capabilities in the context of virtual greeting encounters with their users.

We adopted an interdisciplinary approach that combines multiple background theories in human social psychology for analyzing users' first impressions of an agent based on a series of empirical studies. Then we proposed a new FML specification to represent communicative functions in multimodal communication.

Finally, we designed and implemented a SAIBA compliant computational solution that automates the generation of an agent's nonverbal communicative behavior in virtual greeting encounters in real-time and allows relational agents to manage impressions of personality and attitude towards their users. This solution was designed on the basis of the theoretical foundations provided by our framework and adopted the newly introduced FML specification to represent greeting communicative functions. As part of this solution we also provided a centralized shared FML processor module for multi-agent systems. This module schedules the FML as input and transforms it into BML adopting a data-driven domain independent approach. Finally, we demonstrated a practical application of this solution in the *"Icelandic Language and Culture Training in Virtual Reykjavik"* project.

7.1 Supported Claims

Recalling our original question about properly selecting an agent's nonverbal behavior to avoid unwanted impressions, we demonstrated that nonverbal immediacy cues of smile, gaze and proxemic behavior allows an agent to manage impressions of personality (extraversion) and interpersonal attitude (affiliation) in first greeting encounters with its users. Specifically, we showed that users quickly form impressions of extraversion when greeting agents' proxemics cues are manipulated and interpersonal attitude judgments are made when agents smile and gaze more at the users (see Section 4.2).

We further analyzed the longer-term impact of these selected immediacy cues when exhibited by our agents. In particular, we demonstrated that users' impressions of an agent's attitude (affiliation) overcome impressions of personality (extraversion) when it comes to deciding how likely and for how long they wish to interact with a relational agent. Thus, a favorable impression of a friendly attitude decreases the likelihood that the users reject the agents. We found that managing impressions of attitudes is more important than personality, in particular, by exhibiting more smiling and gazing behavior towards the user during a first greeting encounter, users were more keen to meet and interact with the agents again (see Section 4.3).

7.2 Contributions

To sum up, this thesis makes contributions to different fields of study:

Field of Human Behavior Modeling and First Impressions

We demonstrated that **first impressions** of **personality** and **interpersonal attitude** based on snap judgments of observed nonverbal behavior are still possible when moving from human-human to human-agent interactions.

We modeled nonverbal immediacy behavior into virtual agents by implementing a novel theory that has its foundations in human social psychology work. In addition to contribute to the relational agents field, this implementation has implications for the human social psychology in the context of first impressions. In particular, our theory suggests to consider the impact that someone's nonverbal behavior exhibited during first greeting encounters might have on people's relational decisions.

We showed that the particular **nonverbal behavior** choice of **smile, gaze and proxemics** immediacy cues when carrying out the agents' greeting communicative functions during the phases of a virtual encounter matters in terms of the user's impression formation of the agent's personality and attitude. We also understood that users' first impressions of a relational agent have an impact on users' **relational decisions** in terms of likelihood and frequency of further virtual encounters.

Field of Relational Agents

We combined theories of human social psychology in a theoretical framework that directly fed into Relational Agents work. RAs can now benefit from our findings and increase their utility by managing impressions on users and avoiding an outright rejection at the very first encounter with them.

We provided a **practical demonstration** of our theoretical framework in the "Icelandic Language and Culture Training in Virtual Reykjavik" project. The relational tutoring agents in this 3D virtual learning environment detect approaching users and plan greeting communicative functions in real-time. The multimodal behavior exhibited by the agents in the environment aids both the accomplishment of the fundamental greeting communicative functions and the first impressions management during a first virtual greeting encounter.

Field of Human-Computer Interaction

Our approach smoothly migrated across **different** human-computer interaction **scenarios**. In particular, we deployed virtual agents displayed on regular LCD monitors in the interpretation of behavior study and in the final application in the *"Virtual Reykjavik"* project, then we used life-sized agents in the behavior impact study, and finally we tested a real setting deployment in a public space with Tinker at Boston Museum of Science (see Section 4.4).

We discovered that the communicative greeting models chosen and the concept of interpersonal distance between user and agent suited all of the scenarios that we evaluated.

Field of Computational Linguistics

We identified the major issues to address when designing a representation language for communicative functions (see Sections 5.2.2 and 5.2.3). Then we proposed such specification language with an **FML specification proposal** that addressed several of the issues presented (see Section 5.3).

This specification separates the representation of communicative functions from contextual information, allows a designer to specify coarse-grained temporal constraints between the functions and supports a shared centralized solution to represent communicative functions in multi-agent systems.

Field of System Engineering

We designed and developed a novel SAIBA compliant **computational solution** for existing RA and ECA systems that supports multi-agent systems and the automation of a set of nonverbal behaviors (smile, gaze and proxemics), including positional and orientational parameters, in order to convey specific impressions of personality and attitudes towards the user during the first greeting encounter (see Chapter 6). This solution consists of **two software modules** (named `FML Greeting`

`Agent` and `FML Service`) that can be easily integrated into other systems running on the Unity3D game engine.

The `FML Service` module consists of a shared centralized data-driven FML processor for multi-agent systems. This SAIBA compliant module is domain independent and allows agents to schedule and transform communicative functions (represented in FML) into multimodal behavior (represented in BML).

The practical usage of our FML specification drove both the design and implementation of this computational solution. First, it involved a theoretical formulation for the challenging scheduling problem of FML chunks. We showed a solution for this problem with the design and implementation of a **novel algorithm** presented in Section 6.5.3. Secondly, we showed a rule based technique that merges the contextual information with the scheduled communicative functions provided as input to the `FML Service` to generate BML.

Some final considerations about the adoption of SAIBA as reference architecture should be discussed. Our FML specification supported the design of ECAs running the full SAIBA pipeline. This makes it relatively easy to share our implementation with other SAIBA compliant ECA systems as discussed earlier. However, abstracting the functions specification in FML does not always assure that the supporting multimodal behavior is planned and realized in the same way across systems.

First, the set of of transformation rules need to be shared across systems to obtain the same behavior plan. Secondly, the overall visual result in the hosting ECA system may change. The execution of such plan (i.e. generated BML) can be affected by discrepancies in the realization manner provided by the hosting behavior realizer component or some loss of information between the SAIBA transformation stages. For example, the timing information specified at the FML level might be subject to information loss (i.e. different timing as intended by a designer) when transformed into BML according to the rules, and can be affected by further distortions during the realization stage.

In general, our recommendation is that SAIBA should not only provide standardized interface languages but also techniques and best practices that enable proper transfer between SAIBA components.

7.3 Limitations

We detailed some of the limitations of our approach separately in the conclusion of each previous chapter. However, we summarize here the major limitations concerning the theoretical stances taken, the design and outcome of the evaluation studies presented in our theoretical framework and, finally, technical and practical limitations mainly concerning the FML specification proposed and the computational solution presented in the previous chapter.

Theoretical Limitations

The major theoretical limitation of our approach concerns the adoption of Kendon's greeting model on top of which we manipulated the nonverbal communicative behavior exhibited by our relational agents to manage their impressions. On one hand, this model represented the heart of our theoretical framework and it proved to be transferable from human-human to human-agent interactions. On the other hand, the model only describes 1-to-1 interaction scenarios, thus we were only able to model single user-agent greeting encounters. This had an impact in all of the stages of our work.

From a theoretical perspective, the design of the evaluation studies presented in our framework was affected by this fundamental assumption. For the experiments conducted in a controlled laboratory setting we did not face particular issues. However, in the real setting experimental design with Tinker, we had to face the issue of cleaning up the data gathered, thus loosing a lot of information, from those visitors arriving in groups and therefore not conforming to the assumptions of the model chosen.

From a practical perspective, we implemented our computational solution allowing an individual agent to react towards a single approaching user as our theoretical framework suggested. Thus we limited the applicability of our solution to such 1-to-1 scenarios without the possibility to consider more complex formations where multiple users interact with an individual agent or, in a full multi-party scenario, where a group of users greets a group of agents.

Another limitation of our framework is the lacking of information to model an agent's nonverbal behavior during farewells while holding the impression management of the desired personality and interpersonal attitude. Furthermore, we didn't consider (on a theoretical level) scenarios where (a) the user aborts an ongoing greeting process with an agent but remains in its proximity and (b) the user manifests willingness to "reopen" such greeting. Modeling these aspects would have required the inclusion of new theories in our framework supporting these communicative processes and indicating the proper nonverbal behavior to exhibit. An analysis of the user's behavior in terms of nonverbal input cues indicating when such scenarios are happening would be required as well. These are interesting further developments of our work that we will discuss in the next chapter.

Evaluation Studies Design and Outcome Limitations

In the behavior interpretation study, described in our theoretical framework, we showed that nonverbal behavior interpretation partially depended on users' own personality (see page 51 in Section 4.2.7). However, in the subsequent study on behavior impact we did not have enough information to conclude whether a concordance (or discordance) effect existed between the subject's own personality and the agent's personality and attitude.

The main outcome (i.e. number of visits) of the behavior impact study was a self-report measure. Therefore, we only had a hypothetical decision of the user about the frequency of subsequent encounters with the agents, that we interpreted as increased predicted outcome value assessed by them when evaluating whether to establish a relationship with the agents or not. However, a behavioral measure, for example assessing whether users actually return for the virtual tours they agreed to participate in with the agents prior to beginning the study, would be more informative, but be much more costly in execution.

The effects of managing first impressions in a public space relational agent have been tested, but due to the lacking of empirical evidence we couldn't conclude whether a friendlier version of Tinker would gain higher number of visitors and yield to longer session of interactions with them compared to versions not managing first impressions at all or exhibiting lower friendliness. However, environmental and technical issues that might have affected the overall study outcome have been identified and suggestions for follow-up studies in a similar fashion have been provided (see page 83 in Section 4.4.6).

Technical and Practical Limitations

The FML specification that we proposed requires community feedback prior to be applicable to other RA and ECA systems. In particular, it is important to check whether the theoretical stances taken to identify the FML document structure, the tags and the attributes that we introduced are good. It would be good to formally verify or run further testing on the scheduling algorithm algorithm presented in Section 6.5.3.

Several issues have not been addressed in our FML specification (see Section 5.4 for more details). In particular, an ontology of the information to include in the contextual information and some newly introduced concepts (e.g. ground state) need to be further developed. The transformation rules that we provided are limited to the *initiation* and *closing* interactional communicative functions, but a complete transformer that is capable of mapping functions belonging to all categories and all tracks into BML is ultimately required.

The major technical limitations of our computational solution are related to the theoretical framework limitations concerning the farewell dynamics and greeting re-opening processes. These are missing, but desirable, features that the implemented modules (in particular the `FML Greeting Agent` module) currently do not include. Although we showed a simplified model for implementing farewell dynamics and the user's abortion of a greeting process (i.e. we modeled as intent of finishing the greeting or aborting it prematurely when the user moves away from the agent), the current implementation only partially supports the selection of the agent's nonverbal behavior when the user engages in a greeting process that was prematurely interrupted.

Finally, the desired effects of the nonverbal behavior exhibited by the virtual agents employing our modules for impression management could be easily undermined by the graphical clarity of the animations showing such behavior. We faced this issue in the *"Virtual Reykjavik"* application where some subtle, but at the same time important, smiling cues where not clearly visible from distance (in the particular 3D environment) when the user was at the beginning of the greeting encounter. We solved this problem tweaking the camera parameters and modifying the character's face model, but a solution that is independent by the particular 3D character model and the environment in which is applied is required.

Some of the limitations discussed in this section represent interesting possible future developments of our work that we will discuss in the next, and last, chapter of this thesis.

"The important thing is to never stop questioning. Curiosity has its own reason for existing."
Albert Einstein (1879 – 1955)

8

Future Work

8.1 Short Term

Improving the Theoretical Framework

Kendon's greeting model provided a solid theoretical background for our framework, but at the same time limited the whole approach to 1-to-1 user-agent interactions. A natural question that arises is whether the same model can be adapted to more complex interaction scenarios consisting of 1-to-many (i.e. a user approaching a group of agents), many-to-1 (i.e. many users approaching a single agent) and the most general many-to-many scenario. We foresee that further investigation is needed to understand how some agent's nonverbal cues, for example gaze behavior, are interpreted by users during interaction with such different configurations.

Further exciting explorations are possible. In fact, once an agent (or a group of agents) has properly targeted the nonverbal behavior towards the right addressee(s), it is still a question how impressions of personality and attitudes can be managed. Do our discoveries for 1-to-1 interactions directly migrate to these new configurations?

Regarding the conclusion of a greeting encounter, it is not trivial to properly select the nonverbal behavior that an agent should exhibit in order to consistently manage its impressions as intended since the beginning of the encounter. Therefore, understating the proper nonverbal communicative behavior to exhibit when exchanging farewells while holding certain impressions, it represents an interesting direction that will contribute to improving our framework encompassing all the phases of a greeting encounter.

Finally, we modeled the greeting encounter between the user and the agent as a sequence of phases that an approaching user traverses prior to beginning the conversation with the agent. However, handling situations where a greeting process is interrupted by the user and then reopened, requires a careful study on how an agent should react to this event (while managing its impressions on the user) and what nonverbal behavior choices are more appropriate. For example, a distant salutation might have already occurred during a greeting that is interrupted by the user. When the user engages again in the greeting, even if the distance requires the agent to exhibit behaviors to accomplish the distant salutation, the particular nonverbal behavior exhibited should be carefully planned accordingly and not merely repeated as in the first time.

Furthermore, not all the greetings happen with users walking towards the agents at the same speed, as we assumed in our controlled experiments. Therefore, evaluation studies taking into account various users' walking speed and the adaptation of the agent behavior accordingly are necessary to obtain a greeting model as close as possible to reality.

Advancing the FML Specification

The FML specification proposed with this thesis needs community feedback as part of an iterative process aimed at validating it and encouraging other researchers to participate in this joint effort by improving the specification with further ideas. We plan to keep working on top of this concrete specification and (1) add the missing tags to represent a wider set of communicative functions, (2) complete the mappings from functions to behavior considering all the functions in our proposal, (3) complete the specification of the *ground state* concept in the *mental state* track and (4) provide a more detailed ontology to describe contextual information.

In addition to allow other researchers to apply and test this FML specification in other systems, we also plan to make the scheduling algorithm available to the community in order to test its practical utility and, more in general, to validate our design choices regarding the the temporal constraints in this FML specification.

Making the Implementation Publicly Available

The ease of integration into other RA and ECA systems is a primary goal of our computational solution for impression management. We aim at making it available to the SAIBA community and, therefore, concretely evaluate the flexibility of our solution when applied to different domains and systems employing an actual BML realizer.

Prior to accomplishing this important step, we plan to face with two implementation issues, in particular, concerning the error management and the farewell dynamics as outlined in Section 6.8.1.

8.2 Long Term

Studying First Impressions Impact in the Long Term

Studying the first impressions phenomenon has not merely served the purpose of understanding how to model an agent's nonverbal behavior during the exact moments when the user interacts with the agent. We also aimed at understand whether these formative moments of a first greeting encounter have an impact on users in the long term regarding their choices of continuing the interaction with the agent and building a relationship with it. In order to fully understand this aspect a longitudinal design is required for following up on evaluation studies. This would allow us to examine the stability and long-term impact of the impressions formed in these brief user-agent encounters. A possible design that takes this aspect into account involves the usage of behavioral measures aimed at assessing whether, for example, users actually interact with our relational agents over time and, thus, for how often.

Relevant to this discussion, exploring how long it takes to overcome an agent's impression of personality/interpersonal attitudes managed in a very first encounter is an interesting future direction. In other words, testing how robust are the first impressions formed by users as suggested by [Kammrath et al., 2007] (in that case for impressions of personality traits in human-human interactions).

Exploring Further Dimensions of First Impressions

This work focused on first impressions of a particular Big 5 personality trait, the agent's extraversion level, and the affiliation dimension of Argyle's model of interpersonal attitudes. Considering other personality traits (e.g. neuroticism) and the status dimension (dominance vs. submission) as interpersonal attitude, while keeping the framework that we introduced, represents another intriguing path of research that could be pursued.

We narrowed down the attention to the three fundamental nonverbal immediacy cues of smile, gaze and proxemic behavior, however how users interpret different nonverbal behavior in the context of greeting encounters in terms of the dimensions we focused on or others can be examined (e.g. gesture and posture [Richmond et al., 2008, p. 197]).

Additional modalities such as auditory channels (e.g. verbal salutations) could be further analyzed and modeled in our framework. In particular, it can be interesting to evaluate the amount of information uptake during the impression formation process across modalities, and if users have a preference between visual (i.e. nonverbal cues) and auditory modalities for forming first impressions similarly to the investigations of [Mehrabian and Ferris, 1967].

For a relational agent, considering that the ultimate goal is to establish a relationship with the users, in addition to personality and interpersonal attitude, it might be important to understand how to manage impressions along different dimensions such as *competence* or *trustworthiness*. Some recent work on modeling interpersonal trust in social robots has been done [Lee et al., 2013], but further research in the context of first greeting encounter is still needed. These two dimensions contribute to effectively accomplish an agent's relational goals as suggested by [Bickmore and Cassell, 2001] in the RAs literature and by [Miller et al., 2007] in the human social psychology literature.

Considering Further User Attributes

As suggested by [Krämer et al., 2010], different users with different attributes (i.e. age, gender and computer literacy) interpret an agent's nonverbal behavior differently. Therefore users' attributes in the design of follow-up evaluation studies cannot be neglected.

We certainly need more work to understand the role of the user's own personality when interpreting an agent's nonverbal behavior and expressing a preference for a particular agent type in terms of behavior exhibited. We may also want to look further into possible gender and cultural differences.

The design of follow up evaluation studies can also be improved by considering certain user characteristics and internal states such as the skills in interpreting observed nonverbal behavior, their mood while participating in the studies or their emotional state. A user's skill in interpreting observed nonverbal behavior can greatly impact the results of an agent perception study such as the behavior interpretation study discussed in this thesis. Furthermore, as shown in [Nalini and Skowronski, 2008, p. 94], the user's internal disposition and mood can affect interpersonal judgments and can act as moderator during interpersonal perception studies.

Exploring Different Human-Computer Interaction Settings

We progressively analyzed agents' nonverbal behavior in different human computer interaction settings by deploying virtual agents on regular monitors in the behavior interpretation study, then life-sized agents in the behavior impact study and, finally, detecting user's proximity in the study at the museum with Tinker. The introduction of widely accessible technologies, such as next-generation virtual reality headsets designed to display immersive 3D environments (e.g. Oculus Rift[1]), opens new possibilities for further evaluation studies on first impressions that can have great impact, in particular, in the gaming industry domain.

In the *Virtual Reykjavik* application, the user navigates the environment in first person perspective. This allowed us to focus exclusively on the agent's nonverbal behavior to animate. However, when moving to other applications adopting different camera perspectives, for example in third person view, where the user's driven character (avatar) is visible, the behavior exhibited by this character should be also considered and animated. In fact, the greeting protocol, defined in our framework, requires both the user and the agent to exhibit certain behaviors during the process (e.g. a user might respond to an agent's distant salutation with an hand wave).

In such scenario similar to [Vilhjálmsson and Cassell, 1998, Vilhjálmsson, 2003, Pedica and Vilhjálmsson, 2010], deciding to which extent the users control their own avatars is not a trivial task and it requires facing with several human computer interaction design choices. These choices involve a trade-off between allowing users to control certain behaviors (e.g. moving around) and automate others without requiring the direct intervention of the user (e.g. reacting to a distant salutation). However, the automatic generation of spontaneous reactions should be carefully planned. If the avatar makes a choice that conflicts with what the user had in mind, reliability could be undermined and the user may be left in an uncomfortable state [Vilhjálmsson and Cassell, 1998].

---

[1] See: `www.oculusvr.com`

Bibliography

[Aggarwal and Traum, 2011] Aggarwal, P. and Traum, D. (2011). The bml sequencer: A tool for authoringmulti-character animations. In *Proceedings of the 11th International Conference on Intelligent Virtual Agents*, IVA'11, pages 428–430, Berlin, Heidelberg. Springer-Verlag. (cited at p. 200).

[Al Moubayed et al., 2012] Al Moubayed, S., Beskow, J., Skantze, G., and Granström, B. (2012). Furhat: A back-projected human-like robot head for multiparty human-machine interaction. In Esposito, A., Esposito, A. M., Vinciarelli, A., Hoffmann, R., and Müller, V. C., editors, *Cognitive Behavioural Systems*, volume 7403 of *Lecture Notes in Computer Science*, pages 114–130. Springer Berlin Heidelberg. (cited at p. 38).

[Ambady et al., 1995] Ambady, N., Hallahan, M., and Rosenthal, R. (1995). On judging and being judged accurately in zero-acquaintance situations. *Journal of Personality and Social Psychology*, 69(3):518–529. (cited at p. 52).

[Ambady and Rosenthal, 1992] Ambady, N. and Rosenthal, R. (1992). Thin slices of expressive behavior as predictors of interpersonal consequences: A meta-analysis. *Psychological Bulletin*, 111(2):256–274. (cited at p. 18).

[Ambady and Rosenthal, 1993] Ambady, N. and Rosenthal, R. (1993). Half a minute: Predicting teacher evaluations from thin slices of nonverbal behavior and physical attractiveness. *Journal of Personality and Social Psychology*, 64(3):431–441. (cited at pp. 10, 12, 13).

[Argyle, 1988] Argyle, M. (1988). *Bodily communication*. Methuen, New York, 2nd ed. edition. (cited at pp. 4, 5, 11, 12, 15, 16, 17, 19, 20, 30, 51, 52, 53, 65, 81).

[Argyle and Cook, 1976] Argyle, M. and Cook, M. (1976). *Gaze & Mutual Gaze*, chapter 5, pages 98–124. Cambridge University Press. (cited at p. 5).

[Argyle et al., 1973] Argyle, M., Ingham, R., Alkema, F., and McCallin, M. (1973). The different functions of gaze. *Semiotica*, 7(1):19–32. (cited at p. 5).

[Ballin et al., 2004] Ballin, D., Gillies, M., and Crabtree, B. (2004). A framework for interpersonal attitude and non-verbal communication in improvisational visual media production. In *First European Conference on Visual Media Production IEEE*. (cited at pp. 30, 31).

[Bar et al., 2006] Bar, M., Neta, M., and Linz, H. (2006). Very first impressions. *Emotion*, 6(2):269–278. (cited at p. 12).

[Berger and Kellerman, 1994] Berger, C. and Kellerman, K. (1994). *Strategic Interpersonal Communication*, chapter Acquiring Social Information, pages 1–31. Lawrence Erlbaum Associates, Hillsdale. (cited at p. 11).

[Bergmann et al., 2012] Bergmann, K., Eyssel, F., and Kopp, S. (2012). A second chance to make a first impression? how appearance and nonverbal behavior affect perceived warmth and competence of virtual agents over time. In *Proceedings of the 12th international conference on Intelligent virtual agents*, IVA'12, pages 126–138, Berlin, Heidelberg. Springer-Verlag. (cited at p. 29).

[Bevacqua et al., 2009] Bevacqua, E., Prepin, K., de Sevin, E., Niewiadomski, R., and Pelachaud, C. (2009). Reactive behaviors in saiba architecture. In *Proceedings of the 2nd Function Markup Language Workshop "Towards a Standard markup Langauge for Embodied Dialogue Acts" at AAMAS*. (cited at pp. 102, 138).

[Bickmore, 2008] Bickmore, T. (2008). Framing and interpersonal stance in relational agents. In *Why Conversational Agents do what they do. Workshop on Functional Representations for Generating Conversational Agent Behavior at AAMAS*. (cited at p. 99).

[Bickmore and Cassell, 2001] Bickmore, T. and Cassell, J. (2001). Relational agents: a model and implementation of building user trust. In *Proceedings of SIGCHI*, pages 396–403, New York. ACM. (cited at pp. 6, 14, 32, 217).

[Bickmore et al., 2011] Bickmore, T., Pfeifer, L., and Schulman, D. (2011). Relational agents improve engagement and learning in science museum visitors. In *Proceedings of the 11th international conference on Intelligent virtual agents*, pages 55–67, Berlin, Heidelberg. Springer-Verlag. (cited at p. 79).

[Bickmore et al., 2013] Bickmore, T. W., Pfeifer Vardoulakis, L. M., and Schulman, D. (2013). Tinker: a relational agent museum guide. *Journal of Autonomous*

Agents and Multi-Agent Systems, 27(2):254–276. (cited at pp. 6, 34, 36, 66, 67, 68, 69).

[Bickmore et al., 2005] Bickmore, T. W., Rosalind, and Picard, W. (2005). Establishing and maintaining long-term human-computer relationships. *ACM Transactions on Computer Human Interaction*, 12(2):293–327. (cited at p. 2).

[Bland and Altman G., 1996] Bland, M. J. and Altman G., D. (1996). Statistics notes: Transforming data. *BMJ*, 312(7033):770. (cited at p. 76).

[Bordens and Abbott, 2002] Bordens, K. S. and Abbott, B. B. (2002). *Research design and methods: A process approach*. McGraw-Hill, New York, 8th ed. edition. (cited at p. 58).

[Bradley, 1958] Bradley, J. V. (1958). Complete counterbalancing of immediate sequential effects in a latin square design. *Journal of the American Statistical Association*, 53(282):525–528. (cited at p. 45).

[Brucker, 2007] Brucker, P. (2007). *Scheduling Algorithms*. Springer-Verlag New York, Inc., 5th edition. (cited at pp. 159, 160).

[Burgoon and Le Poire, 1999] Burgoon, J. and Le Poire, B. (1999). Nonverbal cues and interpersonal judgments: Participant and observer perceptions of intimacy, dominance, composure, and formality. *Communication Monographs*, 66(2):105–124. (cited at pp. 14, 17).

[Burgoon et al., 1984] Burgoon, J. K., Buller, D. B., Hale, J. L., and de Turck, M. A. (1984). Relational messages associated with nonverbal behaviors. *Human Communication Research*, 10(3):351–378. (cited at pp. 12, 14, 16, 17, 52, 60).

[Cafaro et al., 2013] Cafaro, A., Vilhjálmsson, H. H., Bickmore, T. W., Heylen, D., and Schulman, D. (2013). First impressions in user-agent encounters: the impact of an agent's nonverbal behavior on users' relational decisions. In *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems*, pages 1201–1202. IFAAMAS. (cited at p. 41).

[Cafaro et al., 2012] Cafaro, A., Vilhjálmsson, H. H., Bickmore, T., Heylen, D., Jóhannsdóttir, K., and Valgarðsson, G. (2012). First impressions: users' judgments of virtual agents' personality and interpersonal attitude in first encounters. In *Proceedings of the 12th international conference on Intelligent Virtual Agents*, IVA'12, pages 67–80. Springer-Verlag. (cited at p. 41).

[Campbell and Rushton, 1978] Campbell, A. and Rushton, J. P. (1978). Bodily communication and personality. *British Journal of Social & Clinical Psychology*, 17(1):31–36. (cited at pp. 12, 16, 51).

[Carifio and Perla, 2007] Carifio, J. and Perla, R. J. (2007). Ten common misunderstandings, misconceptions, persistent myths and urban legends about likert scales and likert response formats and their antidotes. *Journal of Social Sciences*, 3(3):106–116. (cited at p. 47).

[Cassell et al., 1999] Cassell, J., Bickmore, T., Billinghurst, M., Campbell, L., Chang, K., Vilhjálmsson, H., and Yan, H. (1999). Embodiment in conversational interfaces: Rea. In *Proceedings of the SIGCHI conference on Human Factors in Computing Systems*, CHI '99, pages 520–527, New York, NY, USA. ACM. (cited at pp. 1, 87, 90).

[Cassell et al., 2000a] Cassell, J., Bickmore, T., Vilhjálmsson, H., and Yan, H. (2000a). More than just a pretty face: affordances of embodiment. In *Proceedings of the 5th international conference on Intelligent user interfaces*, IUI '00, pages 52–59, New York, NY, USA. ACM. (cited at pp. 34, 36, 93).

[Cassell et al., 2002] Cassell, J., Stocky, T., Bickmore, T., Gao, Y., Nakano, Y., Ryokai, K., Tversky, D., Vaucelle, C., and Vilhjálmsson, H. H. (2002). Mack: Media lab autonomous conversational kiosk. In *Proceedings of Imagina: Intelligent Autonomous Agents*. (cited at pp. 34, 36, 37).

[Cassell et al., 2000b] Cassell, J., Sullivan, J., Prevost, S., and Churchill, E. (2000b). *Embodied Conversational Agents*. The MIT Press, Cambridge, 1st ed. edition. (cited at p. 1).

[Cassell et al., 2001] Cassell, J., Vilhjálmsson, H. H., and Bickmore, T. (2001). Beat: the behavior expression animation toolkit. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, SIGGRAPH '01, pages 477–486, New York, NY, USA. ACM. (cited at pp. 87, 90, 94, 137).

[Cormen et al., 2009] Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2009). *Introduction to Algorithms, Third Edition*. The MIT Press, 3rd edition. (cited at p. 161).

[D'Agostino, 1970] D'Agostino, R. B. (1970). Transformation to normality of the null distribution of g1. *Biometrika*, 57(3):679–681. (cited at p. 77).

[DePaulo, 1992] DePaulo, B. M. (1992). Nonverbal behavior and self-presentation. *Psychological Bulletin*, 111(2):203–243. (cited at pp. 12, 65).

[Deshmukh et al., 2010] Deshmukh, A., Castellano, G., Lim, M. Y., Aylett, R., and McOwan, P. W. (2010). Ubiquitous social perception abilities for interaction initiation in human-robot interaction. In *Proceedings of the 3rd International Workshop on Affective Interaction in Natural Environments*, AFFINE '10, pages 55–60, New York, NY, USA. ACM. (cited at pp. 35, 36).

[Doce et al., 2010] Doce, T., Dias, J., Prada, R., and Paiva, A. (2010). Creating individual agents through personality traits. In *Proceedings of the 10th international conference on Intelligent virtual agents*, pages 257–264, Berlin, Germany. Springer-Verlag. (cited at pp. 30, 31).

[Edwards et al., ress] Edwards, R. A., Bickmore, T., Jenkins, L., Foley, M., and Manjourides, J. (in press). Use of an interactive computer agent to support breastfeeding. *Maternal and Child Health Journal*, pages 1–8. (cited at p. 2).

[Ekman and Friesen, 1978] Ekman, P. and Friesen, W. (1978). *The Facial Action Coding System: A Technique For The Measurement of Facial Movement*. Consulting Psychologists Press, Palo Alto, CA, USA. (cited at p. 155).

[Eysenck, 1991] Eysenck, H. J. (1991). Dimensions of personality: 16, 5 or 3?–Criteria for a taxonomic paradigm. *Personality and Individual Differences*, 12(8):773–790. (cited at p. 17).

[Fukayama et al., 2002] Fukayama, A., Ohno, T., Mukawa, N., Sawaki, M., and Hagita, N. (2002). Messages embedded in gaze of interface agents — impression management with agent's gaze. In *Proceedings of SIGCHI*, pages 41–48, New York. ACM. (cited at p. 28).

[Gelman and Park, 2009] Gelman, A. and Park, D. (2009). Splitting a predictor at the upper quarter or third and the lower quarter or third. *The American Statistician*, 63(1):1–8. (cited at p. 47).

[Gockley et al., 2005] Gockley, R., Bruce, A., Forlizzi, J., Michalowski, M., Mundell, A., Rosenthal, S., Sellner, B., Simmons, R., Snipes, K., Schultz, A., and Wang, J. (2005). Designing robots for long-term social interaction. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2005*, pages 1338–1343. (cited at pp. 37, 67).

[Goffman, 1959] Goffman, E. (1959). *The Presentation of Self in Everyday Life*. Doubleday, Garden City, New York, 1st ed. edition. (cited at pp. 9, 11, 65).

[Goffman, 1981] Goffman, E. (1981). *Forms of Talk*. University of Pennsylvania Press, Philadelphia, PA, USA. (cited at p. 112).

[Goldberg and Saucier, 1998] Goldberg, L. R. and Saucier, G. (1998). What is beyond the big five? *Journal of Personality*, 66(4):495–524. (cited at pp. 17, 30).

[Greenlees et al., 2005] Greenlees, I., Bradley, A., Holder, T., and Thelwell, R. (2005). The impact of opponents' non-verbal behaviour on the first impressions and outcome expectations of table-tennis players. *Journal of Psychology of Sport and Exercise*, 6(1):103–115. (cited at p. 10).

[Hall, 1966] Hall, E. T. (1966). *The Hidden Dimension*. Doubleday. (cited at pp. 20, 21, 34, 44).

[Hartholt et al., 2009] Hartholt, A., Gratch, J., Weiss, L., and Team, T. G. (2009). At the virtual frontier: Introducing gunslinger, a multi-character, mixed-reality, story-driven experience. In Ruttkay, Z., Kipp, M., Nijholt, A., and Vilhjálmsson, H. H., editors, *Intelligent Virtual Agents*, volume 5773 of *Lecture Notes in Computer Science*, pages 500–501. Springer Berlin Heidelberg. (cited at p. 1).

[Hartholt et al., 2013] Hartholt, A., Traum, D. R., Marsella, S. C., Shapiro, A., Stratou, G., Leuski, A., Morency, L.-P., and Gratch, J. (2013). All together now - introducing the virtual human toolkit. In Aylett, R., Krenn, B., Pelachaud, C., and Shimodaira, H., editors, *Intelligent Virtual Agents*, Lecture Notes in Computer Science, pages 368–381. Springer. (cited at pp. 87, 97, 138).

[Heenan et al., 2013] Heenan, B., Greenberg, S., Aghelmanesh, S., and Sharlin, E. (2013). Designing social greetings and proxemics in human robot interaction. Technical report, Department of Computer Science, University of Calgary, Calgary, AB, Canada. (cited at p. 36).

[Heylen et al., 2008] Heylen, D., Kopp, S., Marsella, S. C., Pelachaud, C., and Vilhjálmsson, H. H. (2008). The next step towards a function markup language. In *Proceedings of the 8th international conference on Intelligent Virtual Agents*, IVA '08, pages 270–280, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 87, 91, 98, 105).

[Hoque et al., 2011] Hoque, M., Morency, L.-P., and Picard, R. W. (2011). Are you friendly or just polite? – analysis of smiles in spontaneous face-to-face interactions. In D'Mello, S., Graesser, A., Schuller, B., and Martin, J.-C., editors, *Affective Computing and Intelligent Interaction*, volume 6974 of *Lecture Notes in Computer Science*, pages 135–144. Springer Berlin Heidelberg. (cited at p. 81).

[Hoque et al., 2013] Hoque, M. E., Courgeon, M., Martin, J.-C., Mutlu, B., and Picard, R. W. (2013). Mach: my automated conversation coach. In *Proceedings of* *the 2013 ACM international joint conference on Pervasive and ubiquitous computing*, UbiComp ’13, pages 697–706, New York, NY, USA. ACM. (cited at pp. 1, 2).

[Isbister and Nass, 2000] Isbister, K. and Nass, C. (2000). Consistency of personality in interactive characters: verbal cues, non-verbal cues, and user characteristics. *International Journal of Human-Computer Studies*, 53(2):251–267. (cited at pp. 4, 32, 33).

[Jamieson, 2004] Jamieson, S. (2004). Likert scales: how to (ab)use them. *Medical Education*, 38(12):1217–1218. (cited at p. 46).

[Johnson et al., 2004] Johnson, W. L., Marsella, S., and Vilhjálmsson, H. H. (2004). The darwars tactical language training system. In *The Interservice/Industry Training, Simulation and Education Conference*, pages 1–11. (cited at p. 1).

[Kammrath et al., 2007] Kammrath, L. K., Ames, D. R., and Scholer, A. A. (2007). Keeping up impressions: Inferential rules for impression change across the big five. *Journal of Experimental Social Psychology*, 43(3):450–457. (cited at pp. 4, 18, 216).

[Kang et al., 2012] Kang, S.-H., Gratch, J., Sidner, C., Artstein, R., Huang, L., and Morency, L.-P. (2012). Towards building a virtual counselor: Modeling nonverbal behavior during intimate self-disclosure. In *11th International Conference on Autonomous Agents and Multiagent Systems*, AAMAS ’12. (cited at pp. 2, 28).

[Kang et al., 2008a] Kang, S.-H., Gratch, J., Wang, N., and Watt, J. H. (2008a). Agreeable people like agreeable virtual humans. In *Proceedings of the 8th international conference on Intelligent Virtual Agents*, pages 253–261, Berlin. Springer-Verlag. (cited at pp. 6, 33).

[Kang et al., 2008b] Kang, S.-H., Gratch, J., Wang, N., and Watt, J. H. (2008b). Does the contingency of agents’ nonverbal feedback affect users’ social anxiety? In *Proceedings of the 7th international conference on Autonomous agents and multiagent systems*, pages 120–127. International Foundation for Autonomous Agents and Multiagent Systems. (cited at p. 33).

[Kendon, 1990] Kendon, A. (1990). *Conducting Interaction: Patterns of Behavior in Focused Encounters (Studies in Interactional Sociolinguistics)*. Cambridge University Press, New York. (cited at pp. 4, 20, 21, 22, 23, 24, 25, 34, 44, 65, 72).

[Kopp et al., 2005] Kopp, S., Gesellensetter, L., Krämer, N. C., and Wachsmuth, I. (2005). A conversational agent as museum guide: design and evaluation of a real-world application. In Themis, P., Jonathan, G., Ruth, A., Daniel, B., Patrick, O., and Thomas, R., editors, *Lecture Notes in Computer Science*, pages 329–343. Springer-Verlag. (cited at pp. 6, 37, 67).

[Kopp et al., 2006] Kopp, S., Krenn, B., Marsella, S., Marshall, A. N., Pelachaud, C., Pirker, H., Thórisson, K. R., and Vilhjálmsson, H. H. (2006). Towards a common framework for multimodal generation: the behavior markup language. In *Proceedings of the 6th international conference on Intelligent Virtual Agents*, IVA'06, pages 205–217, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 87, 91, 92, 96).

[Kopp and Pfeiffer-Leßmann, 2008] Kopp, S. and Pfeiffer-Leßmann, N. (2008). Functions of speaking and acting. In *Why Conversational Agents do what they do. Workshop on Functional Representations for Generating Conversational Agent Behavior at AAMAS*. (cited at p. 99).

[Kranstedt et al., 2002] Kranstedt, A., Kopp, S., and Wachsmuth, I. (2002). Murml: A multimodal utterance representation markup language for conversational agents. (cited at p. 94).

[Krämer et al., 2010] Krämer, N., Hoffmann, L., and Kopp, S. (2010). Know your users! empirical results for tailoring an agent's nonverbal behavior to different user groups. In Allbeck, J., Badler, N., Bickmore, T., Pelachaud, C., and Safonova, A., editors, *Intelligent Virtual Agents*, volume 6356 of *Proceedings of Intelligent Virtual Agents (IVA 2010). Lecture Notes in Computer Science, 6356*, pages 468–474. Springer Berlin-Heidelberg. (cited at pp. 32, 217).

[Kulms et al., 2011] Kulms, P., Krämer, N., Gratch, J., and Kang, S.-H. (2011). It's in their eyes: A study on female and male virtual humans' gaze. In *Proceedings of the 11th international conference on Intelligent Virtual Agents*, volume 6895, pages 80–92, Berlin. Springer, Heidelberg. (cited at p. 53).

[Lane et al., 2011] Lane, H. C., Noren, D., Auerbach, D., Birch, M., and Swartout, W. (2011). Intelligent tutoring goes to the museum in the big city: a pedagogical agent for informal science education. In *Proceedings of the 15th international conference on Artificial intelligence in education*, pages 155–162, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 38, 67).

[Leary and Kowalski, 1990] Leary, M. R. and Kowalski, R. M. (1990). Impression management: A literature review and two-component model. *Psychological Bulletin*, 107(1):34–47. (cited at pp. 9, 12).

[Lee and Marsella, 2006] Lee, J. and Marsella, S. (2006). Nonverbal behavior generator for embodied conversational agents. In *Proceedings of the 6th international* *conference on Intelligent Virtual Agents*, IVA'06, pages 243–255, Berlin, Heidelberg. Springer-Verlag. (cited at p. 97).

[Lee and Marsella, 2011] Lee, J. and Marsella, S. (2011). Modeling side participants and bystanders: the importance of being a laugh track. In *Proceedings of the 11th international conference on Intelligent virtual agents*, pages 240–247, Berlin. Springer, Heidelberg. (cited at pp. 30, 31).

[Lee et al., 2013] Lee, J. J., Knox, B., Baumann, J., Breazeal, C., and DeSteno, D. (2013). Computationally modeling interpersonal trust. *Frontiers in Psychology*, 4(893). (cited at p. 217).

[Levesque and Kenny, 1993] Levesque, M. J. and Kenny, D. A. (1993). Accuracy of behavioral predictions at zero acquaintance: A social relations analysis. *Journal of Personality and Social Psychology*, 65(6):1178–1187. (cited at pp. 12, 18).

[Maat and Heylen, 2009] Maat, M. and Heylen, D. (2009). Turn management or impression management? In *Proceedings of the 9th International Conference on Intelligent Virtual Agents*, pages 467–473, Berlin. Springer-Verlag. (cited at p. 28).

[Maat et al., 2010] Maat, M. T., Truong, K. P., and Heylen, D. (2010). How turn-taking strategies influence users' impressions of an agent. In *Proceedings of the 10th International Conference on Intelligent Virtual Agents*, pages 441–453, Berlin. Springer-Verlag. (cited at p. 28).

[MacCallum et al., 2002] MacCallum, R., Zhang, S., Preacher, K., and Rucker, D. (2002). On the practice of dichotomization of quantitative variables. *Psychological Methods*, 7(1):19–40. (cited at p. 47).

[MacLin and Malpass, 2001] MacLin, O. H. and Malpass, R. S. (2001). Racial categorization of faces: The ambiguous race face effect. *Psychology, Public Policy, and Law*, 7(1):98–118. (cited at p. 12).

[Mancini and Pelachaud, 2008] Mancini, M. and Pelachaud, C. (2008). The fml-apml language. In *Why Conversational Agents do what they do. Workshop on Functional Representations for Generating Conversational Agents Behavior at AAMAS*. (cited at pp. 96, 101, 103, 120, 140, 141).

[Marsh and Gilmour, 1988] Marsh, P. E. and Gilmour, R. (1988). *Eye to eye : how people interact*. Salem House Publishers, Topsfield, MA, USA. (cited at pp. 11, 15).

[McCrae and Costa, 1997] McCrae, R. R. and Costa, Paul T., J. (1997). Personality trait structure as a human universal. *American Psychologist*, 52(5):509–516. (cited at p. 17).

[McNeill, 1992] McNeill, D. (1992). *Hand and mind : what gestures reveal about thought*. University of Chicago Press, Chicago. (cited at p. 5).

[Mehrabian, 1969] Mehrabian, A. (1969). Significance of posture and position in the communication of attitude and status relationships. *Psychological Bulletin*, 71(5):359–372. (cited at pp. 12, 15, 16).

[Mehrabian and Ferris, 1967] Mehrabian, A. and Ferris, S. R. (1967). Inference of attitudes from nonverbal communication in two channels. *Consulting Psychology*, 31(3):248–252. (cited at p. 217).

[Mehu et al., 2008] Mehu, M., Little, A. C., and Dunbar, R. I. (2008). Sex differences in the effect of smiling on social judgments: An evolutionary approach. *Journal of Social, Evolutionary, and Cultural Psychology*, 2(3):103–121. (cited at pp. 15, 17, 51, 65, 81).

[Miller et al., 2007] Miller, R., Perlman, D., and Brehm, S. (2007). *Intimate Relationships*. McGraw-Hill, Boston, 4th ed. edition. (cited at pp. 9, 11, 12, 15, 17, 18, 21, 53, 217).

[Moffat, 1997] Moffat, D. (1997). Personality parameters and programs. In Trappl, R. and Petta, P., editors, *Creating Personalities for Synthetic Actors*, volume 1195 of *Lecture Notes in Computer Science*, pages 120–165. Springer-Heidelberg. (cited at p. 18).

[Mohler et al., 2010] Mohler, B. J., Bülthoff, H. H., and Dodds, T. J. (2010). A communication task in hmd virtual environments: Speaker and listener movement improves communication. In *Proceedings of the 23rd Annual Conference on Computer Animation and Social Agents*, pages 1–4. (cited at p. 53).

[Mohler et al., 2008] Mohler, B. J., Bülthoff, H. H., Thompson, W. B., and Creem-Regehr, S. H. (2008). A full-body avatar improves egocentric distance judgments in an immersive virtual environment. In *Proceedings of the 5th symposium on Applied perception in graphics and visualization*, page 194, New York. ACM. (cited at p. 53).

[Moore, 2006] Moore, C. D. (2006). Impression formation. In Ritzer, G., editor, *The Blackwell Encyclopedia of Sociology*, pages 2269–2274. Blackwell Publishing, Malden, MA, USA. (cited at p. 10).

[Nalini and Skowronski, 2008] Nalini, A. and Skowronski, J. J. (2008). *First Impressions*. Guilford Press, New York. (cited at pp. 2, 4, 6, 11, 16, 18, 217).

[Nass et al., 2000] Nass, C., Isbister, K. ., and Lee, E.-J. (2000). Embodied conversational agents. chapter Truth is beauty: researching embodied conversational agents, pages 374–402. MIT Press, Cambridge, MA, USA. (cited at p. 6).

[Nass et al., 1995] Nass, C., Moon, Y., Fogg, B. J., Reeves, B., and Dryer, C. (1995). Can computer personalities be human personalities? In *Conference companion on Human factors in computing systems*, CHI '95, pages 228–229, New York, NY, USA. ACM. (cited at pp. 32, 33).

[Naumann et al., 2009] Naumann, L., Vazire, S., Rentfrow, P., and Gosling, S. (2009). Personality judgments based on physical appearance. *Personality and Social Psychology Bulletin*, 35(12):1661–1671. (cited at p. 12).

[Neff et al., 2011] Neff, M., Toothman, N., Bowmani, R., Tree, J. E. F., and Walker, M. A. (2011). Don't scratch! self-adaptors reflect emotional stability. In *Proceedings of the 11th international conference on Intelligent Virtual Agents*, pages 398–411, Berlin. Springer-Verlag. (cited at pp. 30, 31).

[Neff et al., 2010] Neff, M., Wang, Y., Abbott, R., and Walker, M. (2010). Evaluating the effect of gesture and language on personality perception in conversational agents. In *Proceedings of the 10th international conference on Intelligent Virtual Agents*, pages 222–235, Berlin. Springer-Verlag. (cited at pp. 30, 31).

[Niewiadomski et al., 2009] Niewiadomski, R., Bevacqua, E., Mancini, M., and Pelachaud, C. (2009). Greta: an interactive expressive eca system. In *Proceedings of The 8th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '09, pages 1399–1400. International Foundation for Autonomous Agents and Multiagent Systems. (cited at pp. 87, 90, 96).

[Ochs et al., 2010] Ochs, M., Niewiadomski, R., and Pelachaud, C. (2010). How a virtual agent should smile? In Allbeck, J., Badler, N., Bickmore, T., Pelachaud, C., and Safonova, A., editors, *Intelligent Virtual Agents*, volume 6356 of *Lecture Notes in Computer Science*, pages 427–440. Springer Berlin Heidelberg. (cited at p. 81).

[Pedica and Vilhjálmsson, 2010] Pedica, C. and Vilhjálmsson, H. H. (2010). Spontaneous avatar behavior for human territoriality. *Applied Artificial Intelligence*, 24(6):575–593. (cited at p. 218).

[Poggi, 2007] Poggi, I. (2007). *Mind, hands, face and body: a goal and belief view of multimodal communication*. Weidler, Berlin. (cited at pp. 96, 97).

[Pütten et al., 2010] Pütten, A. M. V. D., Krämer, N. C., and Gratch, J. (2010). How our personality shapes our interactions with virtual characters: implications for research and development. In *Proceedings of the 10th international conference on Intelligent Virtual Agents*, pages 208–221, Berlin. Springer-Verlag. (cited at pp. 6, 33).

[Ravenet et al., 2013] Ravenet, B., Ochs, M., and Pelachaud, C. (2013). From a user-created corpus of virtual agent's non-verbal behavior to a computational model of interpersonal attitudes. In Aylett, R., Krenn, B., Pelachaud, C., and Shimodaira, H., editors, *Intelligent Virtual Agents*, volume 8108 of *Lecture Notes in Computer Science*, pages 263–274. Springer Berlin Heidelberg. (cited at pp. 30, 31).

[Reidsma et al., 2011] Reidsma, D., van Welbergen, H., and Zwiers, J. (2011). Multimodal plan representation for adaptable bml scheduling. In *Proceedings of the 11th international conference on Intelligent virtual agents*, IVA'11, pages 296–308, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 140, 141, 199, 200).

[Richmond et al., 2008] Richmond, V., McCroskey, J., and Hickson, M. (2008). *Nonverbal Communication in Interpersonal Relations*. Allyn and Bacon, 6th ed. edition. (cited at pp. 4, 5, 11, 15, 17, 52, 53, 60, 65, 216).

[Riggio and Friedman, 1986] Riggio, R. E. and Friedman, H. S. (1986). Impression formation: The role of expressive behavior. *Journal of Personality and Social Psychology*, 50(2):421–427. (cited at pp. 11, 12, 52).

[Ring et al., 2013] Ring, L., Barry, B., Totzke, K., and Bickmore, T. (2013). Addressing loneliness and isolation in older adults. In *Affective Computing and Intelligent Interaction*, ACII '13. (cited at pp. 35, 36).

[Rizzo et al., 2013] Rizzo, A., Difede, J., Rothbaum, B. O., Daughtry, J. M., and Reger, G. (2013). Virtual reality as a tool for delivering PTSD exposure therapy. In *Post-Traumatic Stress Disorder: Future Directions in Prevention, Diagnosis, and Treatment*. Springer. (cited at p. 2).

[Rule and Ambady, 2010a] Rule, N. and Ambady, N. (2010a). First impressions of the face: Predicting success. *Social and Personality Psychology Compass*, 4(8):506–516. (cited at p. 11).

[Rule and Ambady, 2010b] Rule, N. and Ambady, N. (2010b). First impressions of the face: Predicting success. *Social and Personality Psychology Compass*, 4(8):506–516. (cited at p. 12).

[Rule and Ambady, 2008] Rule, N. O. and Ambady, N. (2008). Brief exposures: Male sexual orientation is accurately perceived at 50 ms. *Journal of Experimental Social Psychology*, 44(4):1100–1105. (cited at p. 12).

[Rule and Ambady, 2010c] Rule, N. O. and Ambady, N. (2010c). Democrats and republicans can be differentiated from their faces. *PLoS ONE*, 5(1):e8733. (cited at p. 12).

[Rule and Ambady, 2011] Rule, N. O. and Ambady, N. (2011). Judgments of power from college yearbook photos and later career success. *Social Psychological and Personality Science*, 2(154). (cited at p. 12).

[Rule et al., 2009] Rule, N. O., Ambady, N., and Hallett, K. C. (2009). Female sexual orientation is perceived accurately, rapidly, and automatically from the face and its features. *Journal of Experimental Social Psychology*, 45(6):1245–1251. (cited at p. 12).

[Rule et al., 2010] Rule, N. O., Garrett, J. V., and Ambady, N. (2010). On the perception of religious group membership from faces. *PLoS ONE*, 5(12):e14241. (cited at p. 12).

[Salamin et al., 2006] Salamin, P., Thalmann, D., and Vexo, F. (2006). The benefits of third-person perspective in virtual and augmented reality? In *Proceedings of the ACM symposium on Virtual reality software and technology*, pages 27–30, New York. ACM. (cited at p. 53).

[Samtani et al., 2008] Samtani, P., Valente, A., and Johnson, W. L. (2008). Applying the saiba framework to the tactical language and culture training system. In *Why Conversational Agents do what they do. Workshop on Functional Representations for Generating Conversational Agent Behavior at AAMAS*. (cited at pp. 95, 138).

[Saucier, 1994] Saucier, G. (1994). Mini-markers: A brief version of goldberg's unipolar big-five markers. *Journal of Personality Assessment*, 63(3):506. (cited at pp. 45, 46, 58).

[Scherer et al., 2012] Scherer, S., Marsella, S., Stratou, G., Xu, Y., Morbini, F., Egan, A., Rizzo, A., and Morency, L.-P. (2012). Perception markup language: Towards a standardized representation of perceived nonverbal behaviors. In *Intelligent*

Virtual Agents, volume 7502 of *Lecture Notes in Computer Science*, pages 455–463. Springer Berlin Heidelberg. (cited at p. 98).

[Schulman and Bickmore, 2012] Schulman, D. and Bickmore, T. W. (2012). Changes in verbal and nonverbal conversational behavior in long-term interaction. In *Proceedings of the 14th ACM International Conference on Multimodal Interaction*, ICMI '12, pages 11–18, New York, NY, USA. ACM. (cited at p. 29).

[Sevin et al., 2010] Sevin, E. D., Hyniewska, S. J., and Pelachaud, C. (2010). Influence of personality traits on backchannel selection. In *Proceedings of the 10th international conference on Intelligent Virtual Agents*, pages 187–193, Berlin. Springer-Verlag. (cited at pp. 30, 31).

[Sowa et al., 2001] Sowa, T., Latoschik, M., and Kopp, S. (2001). *A Communicative Mediator in a Virtual Environment: Processing of Multimodal Input and Output*, pages 71–74. Proceedings of the International Workshop on Multimodal Presentation and Natural Multimodal Dialogue - IPNMD. (cited at p. 94).

[Stevens and Kristof, 1995] Stevens, C. K. and Kristof, A. L. (1995). Making the right impression: A field study of applicant impression management during job interviews. *Journal of Applied Psychology*, 80(5):587–606. (cited at pp. 10, 12).

[Sunnafrank, 1986] Sunnafrank, M. (1986). Predicted outcome value during initial interactions a reformulation of uncertainty reduction theory. *Human Communication Research*, 13(1):3–33. (cited at pp. 11, 64).

[Sunnafrank and Ramirez, 2004] Sunnafrank, M. and Ramirez, A. (2004). At first sight: Persistent relational effects of get-acquainted conversations. *Journal of Social and Personal Relationships*, 21(3):361–379. (cited at pp. 11, 64).

[Takashima et al., 2008] Takashima, K., Omori, Y., Yoshimoto, Y., Itoh, Y., Kitamura, Y., and Kishino, F. (2008). Effects of avatar's blinking animation on person impressions. In *Proceedings of graphics interface 2008*, pages 169–176, Toronto. Canadian Information Processing Society. (cited at p. 28).

[Tapus et al., 2008] Tapus, A., Ţăpuş, C., and Matarić, M. (2008). User—robot personality matching and assistive robot behavior adaptation for post-stroke rehabilitation therapy. *Intelligent Service Robotics*, 1(2):169–183. (cited at p. 32).

[Thiebaux et al., 2008] Thiebaux, M., Marsella, S., Marshall, A. N., and Kallmann, M. (2008). Smartbody: behavior realization for embodied conversational agents. In *Proceedings of the 7th international joint conference on Autonomous agents and multiagent systems - Volume 1*, AAMAS '08, pages 151–158, Richland, SC.

International Foundation for Autonomous Agents and Multiagent Systems. (cited at pp. 140, 141).

[Thomas and Johnston, 1981] Thomas, F. and Johnston, O. (1981). *Disney animation: The illusion of life*. Hyperion. (cited at p. 4).

[Thórisson, 1994] Thórisson, K. R. (1994). Face-to-face communication with computer agents. In *Proceedings of the AAAI Spring Symposium on Believable Agents*, pages 86–90. (cited at p. 195).

[Thórisson, 1997] Thórisson, K. R. (1997). Gandalf: an embodied humanoid capable of real-time multimodal dialogue with people. In *Proceedings of the first international conference on Autonomous agents*, AGENTS '97, pages 536–537, New York, NY, USA. ACM. (cited at p. 93).

[Thórisson and Vilhjálmsson, 2009] Thórisson, K. R. and Vilhjálmsson, H. H. (2009). Functional description of multimodal acts: A proposal. In *Proceedings of the 2nd Function Markup Language Workshop "Towards a Standard markup Langauge for Embodied Dialogue Acts" at AAMAS*. (cited at pp. 100, 101, 102).

[Traum et al., 2012] Traum, D. R., Aggarwal, P., Artstein, R., Foutz, S., Gerten, J., Katsamanis, A., Leuski, A., Noren, D., and Swartout, W. R. (2012). Ada and grace: Direct interaction with museum visitors. In *Proceedings of the 12th international conference on Intelligent virtual agents*, IVA'12, pages 245–251. Springer-Verlag. (cited at pp. 37, 38, 66, 67).

[Vala et al., 2007] Vala, M., Sequeira, P., Paiva, A., and Aylett, R. (2007). Fearnot! demo: a virtual environment with synthetic characters to help bullying. In *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, AAMAS '07, pages 271:1–271:2, New York, NY, USA. ACM. (cited at p. 1).

[van Oijen, 2007] van Oijen, J. (2007). A framework to support the influence of culture on nonverbal behavior generation in embodied conversational agents. Master's thesis, University of Twente, the Netherlands. (cited at pp. 87, 90, 96, 138).

[van Welbergen et al., 2012] van Welbergen, H., Reidsma, D., and Kopp, S. (2012). An incremental multimodal realizer for behavior co-articulation and coordination. In *Proceedings of the 12th International Conference on Intelligent Virtual Agents*, IVA'12, pages 175–188, Berlin, Heidelberg. Springer-Verlag. (cited at p. 200).

[Vilhjálmsson et al., 2007a] Vilhjálmsson, H. H., Cantelmo, N., Cassell, J., E. Chafai, N., Kipp, M., Kopp, S., Mancini, M., Marsella, S., Marshall, A. N., Pelachaud, C., Ruttkay, Z., Thórisson, K. R., Welbergen, H., and Werf, R. J. (2007a). The behavior markup language: Recent developments and challenges. In *Proceedings of the 7th international conference on Intelligent Virtual Agents*, IVA '07, pages 99–111, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 87, 91, 92, 96).

[Vilhjálmsson and Cassell, 1998] Vilhjálmsson, H. H. and Cassell, J. (1998). Bodychat: Autonomous communicative behaviors in avatars. In *Proceedings of the Second International Conference on Autonomous Agents*, AGENTS '98, pages 269–276, New York, NY, USA. ACM. (cited at pp. 34, 36, 218).

[Vilhjálmsson and Marsella, 2005] Vilhjálmsson, H. H. and Marsella, S. C. (2005). Social performance framework. In *Workshop on Modular Construction of Human-Like Intelligence at the 20th National AAAI Conference on Artificial Intelligence*. (cited at pp. 95, 103).

[Vilhjálmsson et al., 2007b] Vilhjálmsson, H. H., Merchant, C., and Samtani, P. (2007b). Social puppets: Towards modular social animation for agents and avatars. In Schuler, D., editor, *Online Communities and Social Computing*, volume 4564 of *Lecture Notes in Computer Science*, pages 192–201. Springer Berlin Heidelberg. (cited at p. 138).

[Vilhjálmsson, 2003] Vilhjálmsson, H. H. (2003). *Avatar Augmented Online Conversation*. PhD thesis, Media Arts and Sciences, Massachusetts Institute of Technology, Media Laboratory, Cambridge, MA. (cited at p. 218).

[Vilhjálmsson, 2005] Vilhjálmsson, H. H. (2005). Augmenting online conversation through automated discourse tagging. In *Proceedings of the Proceedings of the 38th Annual Hawaii International Conference on System Sciences (HICSS'05) - Track 4 - Volume 04*, HICSS '05, pages 109.1–, Washington, DC, USA. IEEE Computer Society. (cited at pp. 87, 90, 94, 137).

[Vilhjálmsson, 2009] Vilhjálmsson, H. H. (2009). Representing communicative function and behavior in multimodal communication. In Esposito, A., Hussain, A., Marinaro, M., and Martone, R., editors, *Multimodal Signals: Cognitive and Algorithmic Issues*, volume 5398 of *Lecture Notes in Computer Science*, pages 47–59. Springer Berlin Heidelberg. (cited at pp. 5, 14, 91, 106, 108, 117, 118).

[Vinayagamoorthy et al., 2006] Vinayagamoorthy, V., Gillies, M., Steed, A., Tanguy, E., Pan, X., Loscos, C., and Slater, M. (2006). Building expression into virtual characters. In *Eurographics Conference State of the Art Report*, pages 21–61. Eurographics Association. (cited at p. 1).

[Wang and Gratch, 2009] Wang, N. and Gratch, J. (2009). Can virtual human build rapport and promote learning? In *Proceedings of the 2009 conference on Artificial Intelligence in Education*, pages 737–739, Amsterdam, The Netherlands. IOS Press. (cited at p. 29).

[Warwick and Vilhjálmsson, 2005] Warwick, W. and Vilhjálmsson, H. H. (2005). Engendering believable communicative behaviors in synthetic entities for tactical language training: An interim report. In *Proceedings of Behavior Representation in Modeling and Simulation*. (cited at pp. 137, 138).

[Watson et al., 2012] Watson, A., Bickmore, T., Cange, A., Kulshreshtha, A., and Kvedar, J. (2012). An internet-based virtual coach to promote physical activity adherence in overweight adults: Randomized controlled trial. *Journal of Medical Internet Research*, 11(1):293–327. (cited at p. 2).

[Weisbuch et al., 2009] Weisbuch, M., Ivcevic, Z., and Ambady, N. (2009). On being liked on the web and in the "real world": Consistency in first impressions across personal webpages and spontaneous behavior. *Journal of Experimental Social Psychology*, 45(3):573–576. (cited at p. 10).

[Wobbrock et al., 2011] Wobbrock, J. O., Findlater, L., Gergle, D., and Higgins, J. J. (2011). The aligned rank transform for nonparametric factorial analyses using only anova procedures. In *Proceedings of the 2011 annual conference on Human factors in computing systems*, CHI '11, pages 143–146, New York, NY, USA. ACM. (cited at pp. 62, 315).

[Zwiers et al., 2011] Zwiers, J., van Welbergen, H., and Reidsma, D. (2011). Continuous interaction within the saiba framework. In *Proceedings of the 11th international conference on Intelligent virtual agents*, IVA'11, pages 324–330, Berlin, Heidelberg. Springer-Verlag. (cited at pp. 140, 141, 200).

Appendices

A

Nonverbal Behavior Interpretation Study Details

A.1 Manipulation Check

The objective of this manipulation check was threefold:

1. Validate the nonverbal cues exhibited by our greeting agents by making sure that users correctly perceived the difference between the two levels of each one;
2. Adjust the user's avatar locomotion speed during the approach towards the agent as the best compromise between a natural walking speed and enough time given to subjects for observing the behavior exhibited;
3. Ensure that the agent's graphics were visually clear.

Setting and Procedure

We had 10 subjects (2 females and 8 males) that observed, in both camera perspectives, approaches towards the agents similarly to the main study. However, we used a simplified version of the 3D environment with a clear background as depicted in Figure A.1. The subject's avatar walked a *distance* of **13.87 meters** to reach the agent in **12.5 seconds**, thus having a *speed* of **1.11 m/s**.

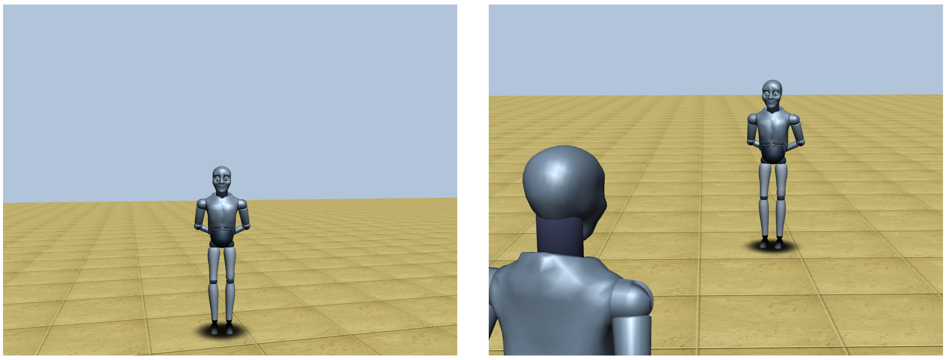

(a) First person perspective (b) Third person perspective

Figure A.1: The simplified deployment of our greeting agent used in the manipulation check both in first (left) and third (right) camera perspective.

Subjects were randomly assigned to an initial camera perspective. The three manipulations (i.e. smile, gaze and proxemics nonverbal cues) were exhibited in isolation from each other. Since there were two levels for each one, subjects observed a succession of two approaches towards the agents for each manipulation (in a fully randomized order), and then replied to a question comparing the pair, as described in the following section. After completing all the approach pairs for a given perspective they repeated the same procedure for the other.

Questions

We used paper and pencil to administer the questionnaires and, given a camera perspective, all the questions were initially hidden and gradually disclosed using separate sheets to avoid priming effects. The full set of questions is listed below. For questions 1-3 there was a follow-up question in case subjects were able to observe a difference (i.e. replying *"Yes"*) between the pair of agents approached (possible answers for these cases are indicated in square brackets). The last two questions were shown after completing all the pairs in a camera perspective mode. They were used to address avatar's speed and graphics related issues (question 4), and to gather general comments and feedbacks (question 5).

1. **Did an agent smile towards you?**
   - No
   - Yes. If so, which one smiled at you? [First Agent | Second Agent | Both]
2. **Did one agent look at you more?**
   - No
   - Yes. Which looked at you more? [First Agent | Second Agent]
3. **Did an agent change its position?**
   - No
   - Yes. If so, how? [It came closer | It moved away]
4. **In general, did you see how the agent reacted towards you?**
   - Yes
   - No, the graphics were visually unclear
   - No, there was not enough time
5. **Any comments?**

Summary of Answers

A summary of subjects' answers grouped by camera perspective is provided in table A.1. The **"No diff."** column represents total number of cases where subjects weren't able to observe any difference. Then we counted the total number of **"Match"** cases. In other words, when in the first approach of a given pair for a manipulation, for example the smile, there was a smiling agent and they correctly marked *"First Agent"*, we counted this as a match. Vice-versa, the **"Mismatch"** column includes cases where subjects attributed the behavior in question to the wrong agent[1]. The total counts in the **"General Question"** raw indicates whether the agent reaction was visible, there were problems with the agent's graphics or there wasn't enough time to observe the reaction.

Cam.	**Manipulation**	**No diff.**	**Match**	**Mismatch**
1P	Smile	2	7	1
	Gaze	1	8	1
	Proxemics	2	7	1
	General Question	**Reaction Visible**	**Graphics Problem**	**Time Problem**
		10	0	0
Cam.	**Manipulation**	**No diff.**	**Match**	**Mismatch**
3P	Smile	2	7	1
	Gaze	3	6	1
	Proxemics	0	10	0
	General Question	**Reaction Visible**	**Graphics Problem**	**Time Problem**
		10	0	0

Table A.1: Summary of answers provided by subjects in the manipulation check grouped by camera perspective (**Cam.**). For the three manipulations, the columns indicate total counts of cases where subjects: didn't see any difference (**No diff.**), correctly matched their answer with the levels observed in succession (**Match**) and, vice-versa, attributed a behavior to the wrong agent in a pair (**Mismatch**). The **General Question** raw reports total counts for question 4.

From these responses we concluded that the number of matches was relatively higher compared to the mismatches, thus subjects were able to observe differences in the manipulations observed. Subjects didn't report any problem related to the graphics and the duration of the approaches, therefore we adopted the same avatar's speed in the main study as well. There weren't major issues associated with their feedbacks.

---

[1] The count of mismatches for question 1 includes also cases where subject replied with "Both".

A.2 Questionnaires

Main and Exploratory Questions

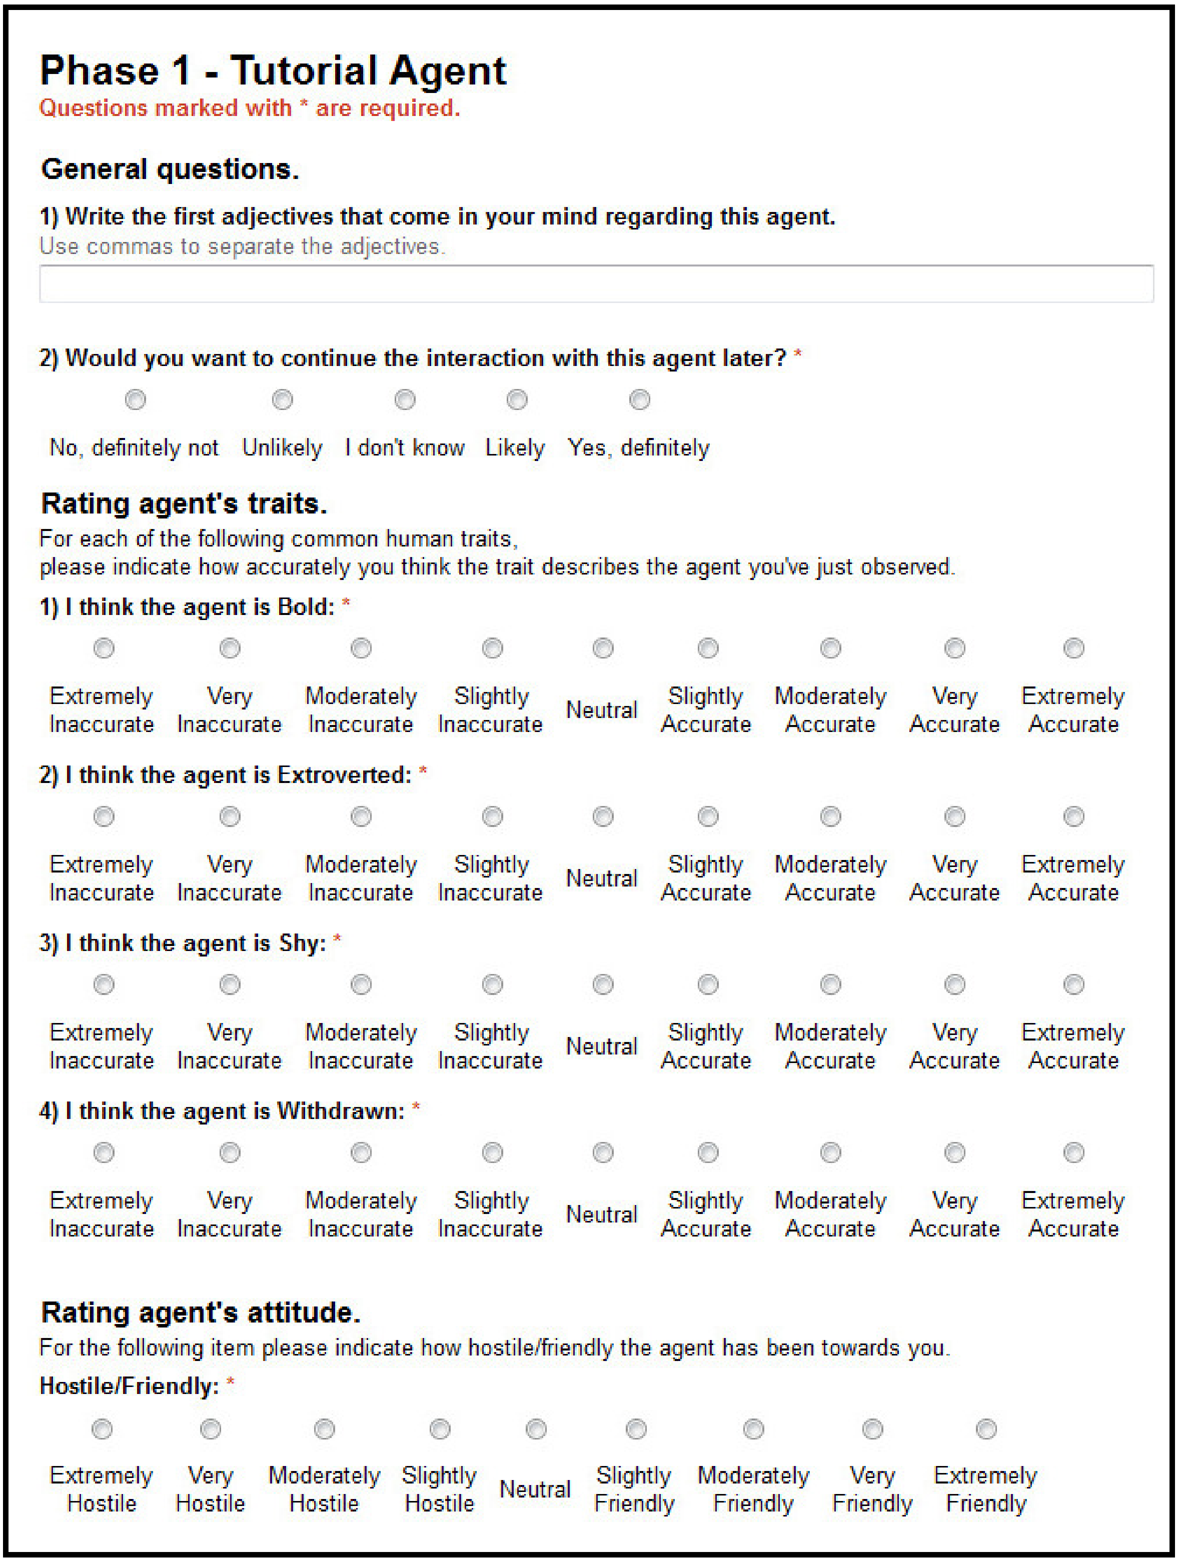

Phase 1 - Tutorial Agent

Questions marked with * are required.

General questions.

1) Write the first adjectives that come in your mind regarding this agent.

Use commas to separate the adjectives.

2) Would you want to continue the interaction with this agent later? *

○	○	○	○	○
No, definitely not	Unlikely	I don't know	Likely	Yes, definitely

Rating agent's traits.

For each of the following common human traits,
please indicate how accurately you think the trait describes the agent you've just observed.

1) I think the agent is Bold: *

○	○	○	○	○	○	○	○	○
Extremely Inaccurate	Very Inaccurate	Moderately Inaccurate	Slightly Inaccurate	Neutral	Slightly Accurate	Moderately Accurate	Very Accurate	Extremely Accurate

2) I think the agent is Extroverted: *

○	○	○	○	○	○	○	○	○
Extremely Inaccurate	Very Inaccurate	Moderately Inaccurate	Slightly Inaccurate	Neutral	Slightly Accurate	Moderately Accurate	Very Accurate	Extremely Accurate

3) I think the agent is Shy: *

○	○	○	○	○	○	○	○	○
Extremely Inaccurate	Very Inaccurate	Moderately Inaccurate	Slightly Inaccurate	Neutral	Slightly Accurate	Moderately Accurate	Very Accurate	Extremely Accurate

4) I think the agent is Withdrawn: *

○	○	○	○	○	○	○	○	○
Extremely Inaccurate	Very Inaccurate	Moderately Inaccurate	Slightly Inaccurate	Neutral	Slightly Accurate	Moderately Accurate	Very Accurate	Extremely Accurate

Rating agent's attitude.

For the following item please indicate how hostile/friendly the agent has been towards you.

Hostile/Friendly: *

○	○	○	○	○	○	○	○	○
Extremely Hostile	Very Hostile	Moderately Hostile	Slightly Hostile	Neutral	Slightly Friendly	Moderately Friendly	Very Friendly	Extremely Friendly

Personality Inventory

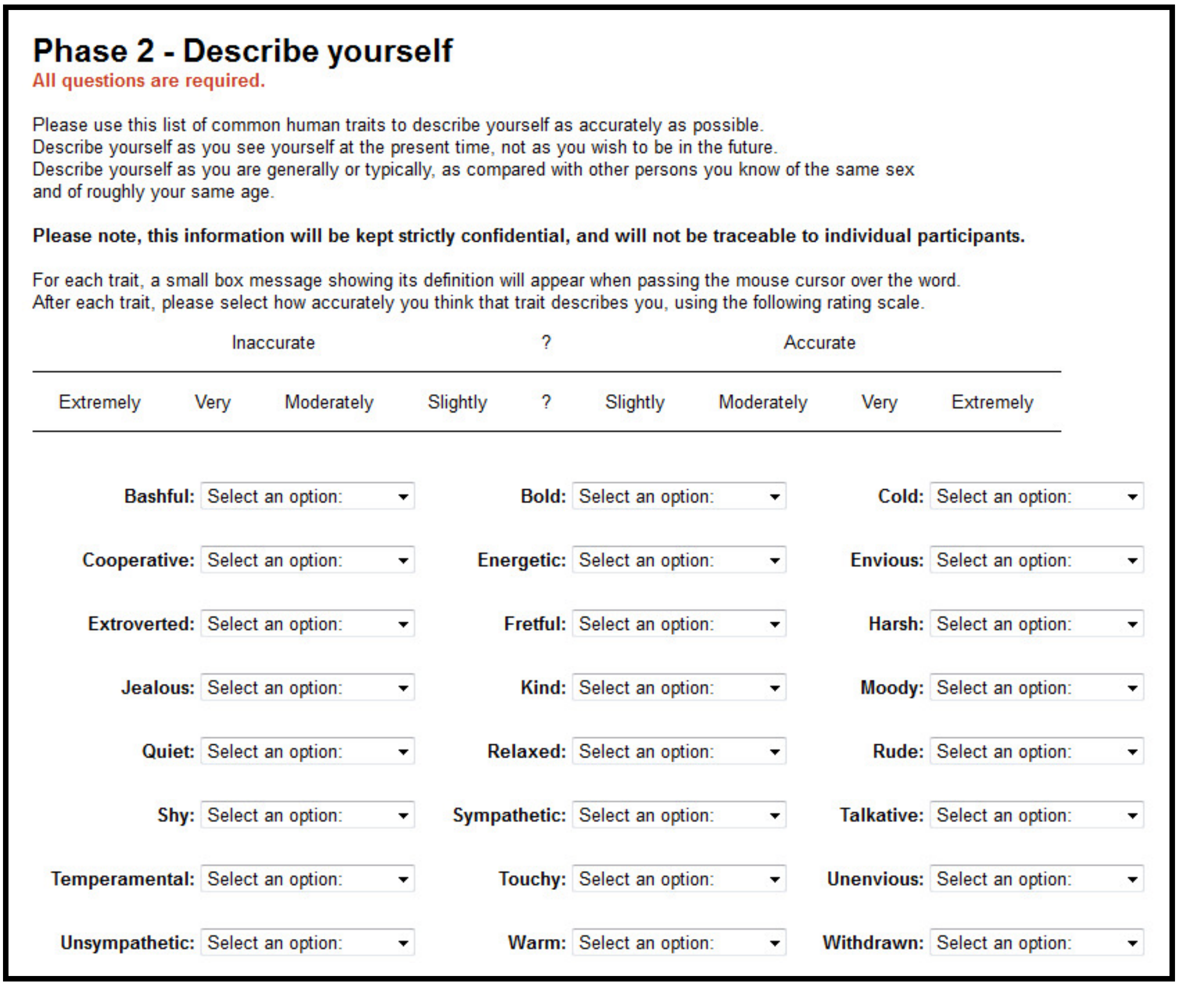

Phase 2 - Describe yourself

All questions are required.

Please use this list of common human traits to describe yourself as accurately as possible.
Describe yourself as you see yourself at the present time, not as you wish to be in the future.
Describe yourself as you are generally or typically, as compared with other persons you know of the same sex and of roughly your same age.

Please note, this information will be kept strictly confidential, and will not be traceable to individual participants.

For each trait, a small box message showing its definition will appear when passing the mouse cursor over the word.
After each trait, please select how accurately you think that trait describes you, using the following rating scale.

Inaccurate				?	Accurate			
Extremely	Very	Moderately	Slightly	?	Slightly	Moderately	Very	Extremely

Bashful:	Select an option: ▾	**Bold:**	Select an option: ▾	**Cold:**	Select an option: ▾
Cooperative:	Select an option: ▾	**Energetic:**	Select an option: ▾	**Envious:**	Select an option: ▾
Extroverted:	Select an option: ▾	**Fretful:**	Select an option: ▾	**Harsh:**	Select an option: ▾
Jealous:	Select an option: ▾	**Kind:**	Select an option: ▾	**Moody:**	Select an option: ▾
Quiet:	Select an option: ▾	**Relaxed:**	Select an option: ▾	**Rude:**	Select an option: ▾
Shy:	Select an option: ▾	**Sympathetic:**	Select an option: ▾	**Talkative:**	Select an option: ▾
Temperamental:	Select an option: ▾	**Touchy:**	Select an option: ▾	**Unenvious:**	Select an option: ▾
Unsympathetic:	Select an option: ▾	**Warm:**	Select an option: ▾	**Withdrawn:**	Select an option: ▾

Demographics

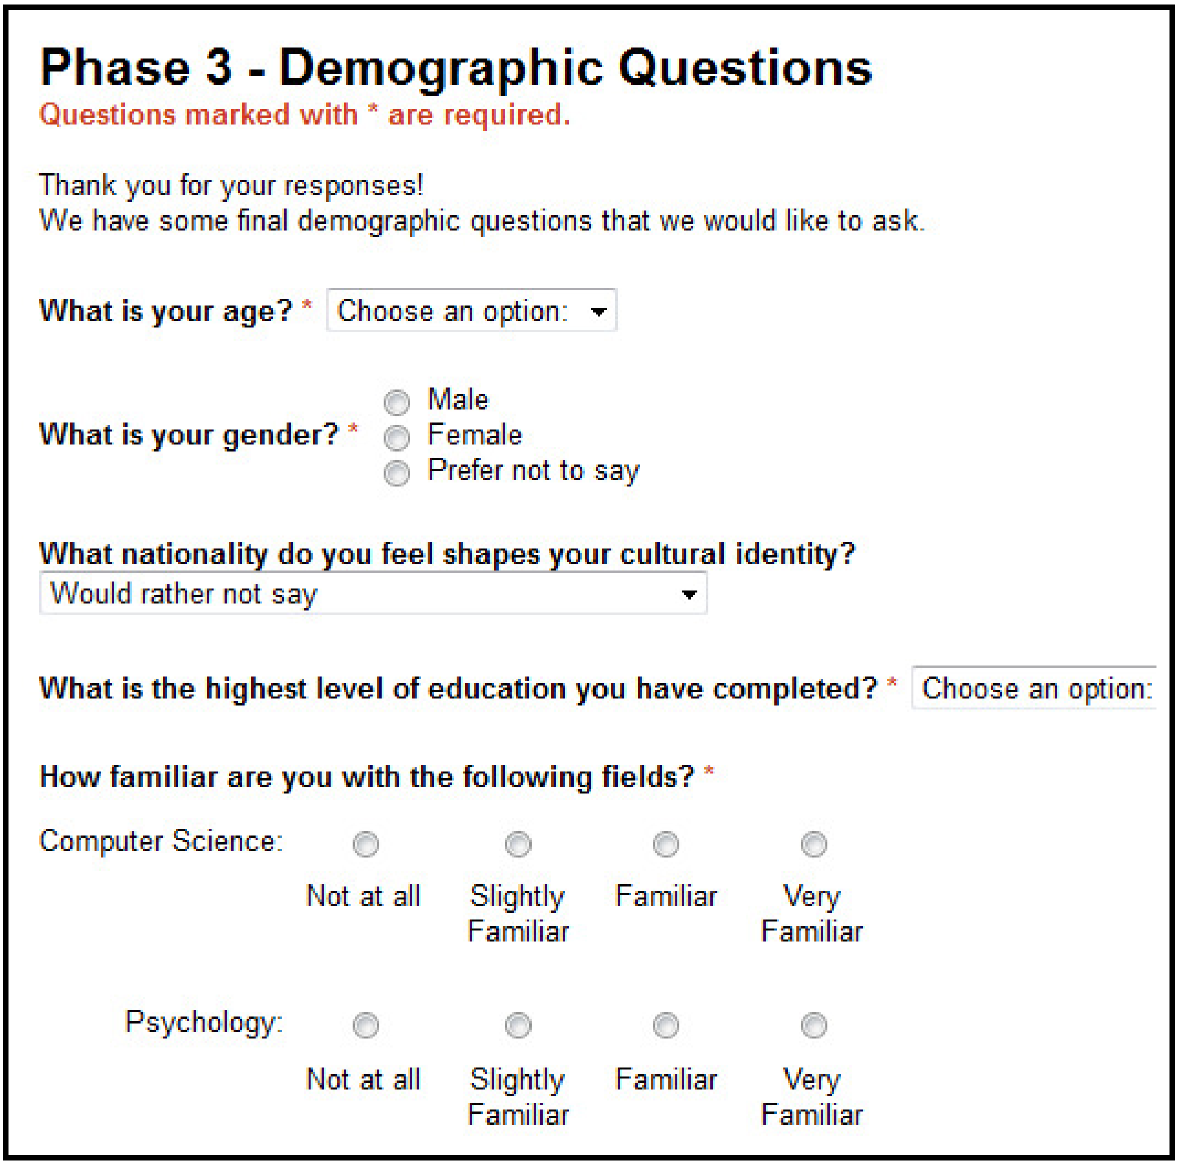

Phase 3 - Demographic Questions

Questions marked with * are required.

Thank you for your responses!
We have some final demographic questions that we would like to ask.

What is your age? * Choose an option:

What is your gender? *
○ Male
○ Female
○ Prefer not to say

What nationality do you feel shapes your cultural identity?
Would rather not say

What is the highest level of education you have completed? * Choose an option:

How familiar are you with the following fields? *

Computer Science:	○ Not at all	○ Slightly Familiar	○ Familiar	○ Very Familiar
Psychology:	○ Not at all	○ Slightly Familiar	○ Familiar	○ Very Familiar

A.3 Demographics

Subjects Age

Range	Frequency 1P Trial	Frequency 3P Trial
< 18	-	-
18-20	-	-
21-30	20	20
31-40	10	5
41-50	-	6
51-60	2	1
> 61	-	32
Total	32	32

Table A.2: Subjects age frequencies in range intervals (shown in years) for the first person view (1P) and third person view (3P) trials.

Subjects Cultural Identity

Country	Frequency 1P Trial	Frequency 3P Trial
Canada	-	1
Czech Republic	-	1
Germany	3	-
Iceland	17	24
India	2	-
Iran	2	-
Italy	2	1
Lebanon	1	-
Lithuania	1	-
Mexico	1	1
Nigeria	1	-
Poland	1	-
Romania	-	1
Spain	-	1
USA	1	1
Not said	-	1
Total	32	32

Table A.3: Subjects cultural identity frequencies by country for the first person view (1P) and third person view (3P) trials.

Subjects Level of Education

Education Level	Frequency 1P Trial	Frequency 3P Trial
Doctorate level	4	2
Master level	8	6
Undergraduate level	11	7
High school	8	16
Elementary school	1	-
Less than elementary school	-	-
Other	-	1
Total	32	32

Table A.4: Subjects level of education frequencies for the first person view (1P) and third person view (3P) trials.

A.4 Documents

Consent Declaration Form

Behavior Interpretation Study

CONSENT DECLARATION
Jan 2012

Principal Investigator:
Angelo Cafaro, CADIA, School of Computer Science, Reykjavik University

In this study, titled "Behavior Interpretation Study", I will be interacting with some graphical agents in the main reception of a virtual museum. The interactions with each agent will be shown on a 19" LCD monitor and realized by a system running on a personal computer. I fully understand that my participation is voluntary and that I am free to withdraw my consent and to discontinue participation at *any time* without prejudice to myself. The experimental procedure has been explained to me and the Investigator, Angelo Cafaro, has offered to answer any inquiries concerning this procedure.

I have the option to leave my e-mail address, by doing so I accept to participate to a lottery involving all subjects participating in this study and I allow the Investigator to contact me by e-mail in case I'll be drawn. The prize for this lottery is a gift card for "Te & Kaffi" and participation is free of any charge.

I understand that I can contact the Director of Research Services at Reykjavik University, Kristján Kristjánsson (kk@ru.is), if I believe I have been treated unfairly as a subject and/or I believe that the research team has breached the RU Code of Ethics.

I have read and understood the above, as well as the experiment instructions, and agree to participate in this research effort.

FULL NAME OF SUBJECT

E-MAIL (OPTIONAL)

SIGNATURE and DATE

Instructions

INSTRUCTIONS: BEHAVIOR INTERPRETATION STUDY

You will be interacting with eight different Intelligent Virtual Agents (IVAs) that have the role of stewards in a virtual museum's main reception. We are going to ask you to observe how every single agent reacts towards you during each approach and then reply to some questions regarding the first impressions you might have gotten about the observed steward.

Intelligent Virtual Agents are interactive characters capable of communicating with users, or among their selves, using human-like natural modalities, such as verbal or non-verbal behaviors. An Avatar (body that represents you) is a graphical representation of the user in a virtual environment. In this study, you have your own avatar, the steward is the IVA and the virtual museum entrance is the virtual environment.

Every approach will involve you with your avatar and a ***single*** steward agent at time. It will consist of the following steps:

1. You will see the environment from a first (third) person perspective, which means from the point of view of your avatar (which means you can see the body of your avatar on-screen from behind);
2. You will start with your avatar standing outside the main entrance. There will be one steward agent inside;
3. You will approach the steward to ask some information regarding the museum. You will be able to start this whenever you are ready by pressing a button;
4. As the approach begins, you will be observing how the steward reacts.
5. When the approach is complete, your avatar will stop automatically and you will be asked some questions regarding the impressions you might have formed of the steward you have just approached. When you have finished answering ***all required*** questions, you can proceed to the next approach pressing a button.
   Please note*: you can't watch a scene again, therefore do your best answering the questions even though you are not sure about the scene you've just observed;*
6. The procedure described above will be the same for all the other steward agents. The total number of approaches you will be observing is **eight (8)**. After you have completed this phase, you will be required to fill in a short **Personality Inventory** about yourself and, finally, you will be asked to provide some **Demographic Information**.

In sum, the whole procedure is composed of the following **main phases**:

1. Observe and reply questions about the specific approach 8 times;
2. Fill in a short personality inventory about yourself;
3. And finally, give some demographic information.

In order for you to get familiar with the first phase, we will complete together a single approach which will serve as training for you. After that I'll let you continue with Phase 1 on your own.

If you have any questions, feel free to ask now. After the training approach we'll leave you to yourself (unless the system crashes, in which case, ask me for help simply waving, I'll be reading in the adjacent room).

Debriefing

Behavior Interpretation Study

DEBRIEFING
Jan 2012

Principal Investigator:
Angelo Cafaro, CADIA, School of Computer Science, Reykjavik University

You have just participated in a study intended to help understanding how people interpret IVAs' non-verbal behaviors when they approach them for the first time with their own avatar. You have been observing a variation of some non-verbal behavior assigned to each steward and rated the impressions each of them made on you according to his extraversion/introversion and hostility/friendliness.

Each steward agent you have been observing was programmed to show you specific Smile, Gaze and Proxemics (movement) behavior. The particular assignment was obtained by picking the following options for each of them:

- Smile: smiling when you approached him vs. not smiling;
- Gaze: high percentage during the approach vs. low;
- Proxemics: a step towards you at the end of the approach vs. staying still the whole time.

The particular questions given after observing each steward's behavior were addressing impressions of his personality (extraverted/introverted) and his attitude towards you (friendly/hostile). In the end, the personality inventory you have filled in, will be used to see if there exists any kind of correlation between the ratings you gave to each agent with your own personality.

Your personal information will be kept strictly confidential, will not be traceable to individual participants and will not be sold, reused, rented, loaned or otherwise disclosed. Any information you give us will not be used in ways that you have not consented to.

If you, for any reason, are not comfortable with the way the experiment was conducted, feel free to withdraw your consent declaration and cancel your participation.

We kindly ask you to keep secret this short explanation of the study purposes, since we will be running this study on other persons you might know or you might have talked about.

A.5 Summary of Means and ANOVA Tables

First Person Perspective Trial

Agent Extraversion Measure

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Smile	Sphericity Assumed	2,409	1	2,409	,937	,342	,036	,937	,154
	Greenhouse-Geisser	2,409	1,000	2,409	,937	,342	,036	,937	,154
	Huynh-Feldt	2,409	1,000	2,409	,937	,342	,036	,937	,154
	Lower-bound	2,409	1,000	2,409	,937	,342	,036	,937	,154
Smile * Subject Extraversion	Sphericity Assumed	10,924	2	5,462	2,125	,140	,145	4,250	,394
	Greenhouse-Geisser	10,924	2,000	5,462	2,125	,140	,145	4,250	,394
	Huynh-Feldt	10,924	2,000	5,462	2,125	,140	,145	4,250	,394
	Lower-bound	10,924	2,000	5,462	2,125	,140	,145	4,250	,394
Smile * Subject Agreeableness	Sphericity Assumed	14,088	2	7,044	2,741	,084	,180	5,482	,492
	Greenhouse-Geisser	14,088	2,000	7,044	2,741	,084	,180	5,482	,492
	Huynh-Feldt	14,088	2,000	7,044	2,741	,084	,180	5,482	,492
	Lower-bound	14,088	2,000	7,044	2,741	,084	,180	5,482	,492
Smile * Subject Neuroticism	Sphericity Assumed	4,305	2	2,152	,838	,445	,063	1,675	,177
	Greenhouse-Geisser	4,305	2,000	2,152	,838	,445	,063	1,675	,177
	Huynh-Feldt	4,305	2,000	2,152	,838	,445	,063	1,675	,177
	Lower-bound	4,305	2,000	2,152	,838	,445	,063	1,675	,177
Error(Smile)	Sphericity Assumed	64,252	25	2,570					
	Greenhouse-Geisser	64,252	25,000	2,570					
	Huynh-Feldt	64,252	25,000	2,570					
	Lower-bound	64,252	25,000	2,570					
Gaze	Sphericity Assumed	5,474	1	5,474	3,281	,082	,116	3,281	,414
	Greenhouse-Geisser	5,474	1,000	5,474	3,281	,082	,116	3,281	,414
	Huynh-Feldt	5,474	1,000	5,474	3,281	,082	,116	3,281	,414
	Lower-bound	5,474	1,000	5,474	3,281	,082	,116	3,281	,414
Gaze * Subject Extraversion	Sphericity Assumed	11,169	2	5,584	3,347	,052	,211	6,693	,579
	Greenhouse-Geisser	11,169	2,000	5,584	3,347	,052	,211	6,693	,579
	Huynh-Feldt	11,169	2,000	5,584	3,347	,052	,211	6,693	,579
	Lower-bound	11,169	2,000	5,584	3,347	,052	,211	6,693	,579
Gaze * Subject Agreeableness	Sphericity Assumed	8,369	2	4,184	2,508	,102	,167	5,015	,456
	Greenhouse-Geisser	8,369	2,000	4,184	2,508	,102	,167	5,015	,456
	Huynh-Feldt	8,369	2,000	4,184	2,508	,102	,167	5,015	,456
	Lower-bound	8,369	2,000	4,184	2,508	,102	,167	5,015	,456
Gaze * Subject Neuroticism	Sphericity Assumed	3,746	2	1,873	1,123	,341	,082	2,245	,225
	Greenhouse-Geisser	3,746	2,000	1,873	1,123	,341	,082	2,245	,225
	Huynh-Feldt	3,746	2,000	1,873	1,123	,341	,082	2,245	,225
	Lower-bound	3,746	2,000	1,873	1,123	,341	,082	2,245	,225
Error(Gaze)	Sphericity Assumed	41,716	25	1,669					
	Greenhouse-Geisser	41,716	25,000	1,669					
	Huynh-Feldt	41,716	25,000	1,669					
	Lower-bound	41,716	25,000	1,669					
Proxemics	Sphericity Assumed	101,811	1	101,811	39,229	,000	,611	39,229	1,000
	Greenhouse-Geisser	101,811	1,000	101,811	39,229	,000	,611	39,229	1,000
	Huynh-Feldt	101,811	1,000	101,811	39,229	,000	,611	39,229	1,000
	Lower-bound	101,811	1,000	101,811	39,229	,000	,611	39,229	1,000
Proxemics * Subject Extraversion	Sphericity Assumed	4,174	2	2,087	,804	,459	,060	1,608	,172
	Greenhouse-Geisser	4,174	2,000	2,087	,804	,459	,060	1,608	,172
	Huynh-Feldt	4,174	2,000	2,087	,804	,459	,060	1,608	,172
	Lower-bound	4,174	2,000	2,087	,804	,459	,060	1,608	,172
Proxemics * Subject Agreeableness	Sphericity Assumed	6,342	2	3,171	1,222	,312	,089	2,444	,242
	Greenhouse-Geisser	6,342	2,000	3,171	1,222	,312	,089	2,444	,242
	Huynh-Feldt	6,342	2,000	3,171	1,222	,312	,089	2,444	,242
	Lower-bound	6,342	2,000	3,171	1,222	,312	,089	2,444	,242

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Proxemics * Subject Neuroticism	Sphericity Assumed	2,646	2	1,323	,510	,607	,039	1,020	,124
	Greenhouse-Geisser	2,646	2,000	1,323	,510	,607	,039	1,020	,124
	Huynh-Feldt	2,646	2,000	1,323	,510	,607	,039	1,020	,124
	Lower-bound	2,646	2,000	1,323	,510	,607	,039	1,020	,124
Error(Proximity)	Sphericity Assumed	64,883	25	2,595					
	Greenhouse-Geisser	64,883	25,000	2,595					
	Huynh-Feldt	64,883	25,000	2,595					
	Lower-bound	64,883	25,000	2,595					
Smile * Gaze	Sphericity Assumed	,971	1	,971	,662	,423	,026	,662	,123
	Greenhouse-Geisser	,971	1,000	,971	,662	,423	,026	,662	,123
	Huynh-Feldt	,971	1,000	,971	,662	,423	,026	,662	,123
	Lower-bound	,971	1,000	,971	,662	,423	,026	,662	,123
Smile * Gaze * Subject Extraversion	Sphericity Assumed	,364	2	,182	,124	,884	,010	,248	,067
	Greenhouse-Geisser	,364	2,000	,182	,124	,884	,010	,248	,067
	Huynh-Feldt	,364	2,000	,182	,124	,884	,010	,248	,067
	Lower-bound	,364	2,000	,182	,124	,884	,010	,248	,067
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	5,061	2	2,530	1,726	,199	,121	3,452	,328
	Greenhouse-Geisser	5,061	2,000	2,530	1,726	,199	,121	3,452	,328
	Huynh-Feldt	5,061	2,000	2,530	1,726	,199	,121	3,452	,328
	Lower-bound	5,061	2,000	2,530	1,726	,199	,121	3,452	,328
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	3,717	2	1,859	1,268	,299	,092	2,536	,250
	Greenhouse-Geisser	3,717	2,000	1,859	1,268	,299	,092	2,536	,250
	Huynh-Feldt	3,717	2,000	1,859	1,268	,299	,092	2,536	,250
	Lower-bound	3,717	2,000	1,859	1,268	,299	,092	2,536	,250
Error(Smile*Gaze)	Sphericity Assumed	36,650	25	1,466					
	Greenhouse-Geisser	36,650	25,000	1,466					
	Huynh-Feldt	36,650	25,000	1,466					
	Lower-bound	36,650	25,000	1,466					
Smile * Proxemics	Sphericity Assumed	2,477	1	2,477	1,085	,308	,042	1,085	,171
	Greenhouse-Geisser	2,477	1,000	2,477	1,085	,308	,042	1,085	,171
	Huynh-Feldt	2,477	1,000	2,477	1,085	,308	,042	1,085	,171
	Lower-bound	2,477	1,000	2,477	1,085	,308	,042	1,085	,171
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	5,802	2	2,901	1,271	,298	,092	2,541	,250
	Greenhouse-Geisser	5,802	2,000	2,901	1,271	,298	,092	2,541	,250
	Huynh-Feldt	5,802	2,000	2,901	1,271	,298	,092	2,541	,250
	Lower-bound	5,802	2,000	2,901	1,271	,298	,092	2,541	,250
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	5,679	2	2,839	1,244	,306	,090	2,487	,245
	Greenhouse-Geisser	5,679	2,000	2,839	1,244	,306	,090	2,487	,245
	Huynh-Feldt	5,679	2,000	2,839	1,244	,306	,090	2,487	,245
	Lower-bound	5,679	2,000	2,839	1,244	,306	,090	2,487	,245
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	1,673	2	,836	,366	,697	,028	,733	,102
	Greenhouse-Geisser	1,673	2,000	,836	,366	,697	,028	,733	,102
	Huynh-Feldt	1,673	2,000	,836	,366	,697	,028	,733	,102
	Lower-bound	1,673	2,000	,836	,366	,697	,028	,733	,102
Error(Smile*Proxemics)	Sphericity Assumed	57,077	25	2,283					
	Greenhouse-Geisser	57,077	25,000	2,283					
	Huynh-Feldt	57,077	25,000	2,283					
	Lower-bound	57,077	25,000	2,283					
Gaze * Proxemics	Sphericity Assumed	,136	1	,136	,067	,797	,003	,067	,057
	Greenhouse-Geisser	,136	1,000	,136	,067	,797	,003	,067	,057
	Huynh-Feldt	,136	1,000	,136	,067	,797	,003	,067	,057
	Lower-bound	,136	1,000	,136	,067	,797	,003	,067	,057
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	6,867	2	3,433	1,707	,202	,120	3,413	,324
	Greenhouse-Geisser	6,867	2,000	3,433	1,707	,202	,120	3,413	,324
	Huynh-Feldt	6,867	2,000	3,433	1,707	,202	,120	3,413	,324
	Lower-bound	6,867	2,000	3,433	1,707	,202	,120	3,413	,324
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	1,731	2	,866	,430	,655	,033	,861	,112
	Greenhouse-Geisser	1,731	2,000	,866	,430	,655	,033	,861	,112
	Huynh-Feldt	1,731	2,000	,866	,430	,655	,033	,861	,112
	Lower-bound	1,731	2,000	,866	,430	,655	,033	,861	,112
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	3,895	2	1,948	,968	,394	,072	1,936	,199
	Greenhouse-Geisser	3,895	2,000	1,948	,968	,394	,072	1,936	,199

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	3,895	2,000	1,948	,968	,394	,072	1,936	,199
	Lower-bound	3,895	2,000	1,948	,968	,394	,072	1,936	,199
Error(Gaze*Proxemics)	Sphericity Assumed	50,291	25	2,012					
	Greenhouse-Geisser	50,291	25,000	2,012					
	Huynh-Feldt	50,291	25,000	2,012					
	Lower-bound	50,291	25,000	2,012					
Smile * Gaze * Proxemics	Sphericity Assumed	,497	1	,497	,232	,634	,009	,232	,075
	Greenhouse-Geisser	,497	1,000	,497	,232	,634	,009	,232	,075
	Huynh-Feldt	,497	1,000	,497	,232	,634	,009	,232	,075
	Lower-bound	,497	1,000	,497	,232	,634	,009	,232	,075
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	7,341	2	3,671	1,718	,200	,121	3,436	,326
	Greenhouse-Geisser	7,341	2,000	3,671	1,718	,200	,121	3,436	,326
	Huynh-Feldt	7,341	2,000	3,671	1,718	,200	,121	3,436	,326
	Lower-bound	7,341	2,000	3,671	1,718	,200	,121	3,436	,326
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	3,262	2	1,631	,763	,477	,058	1,527	,165
	Greenhouse-Geisser	3,262	2,000	1,631	,763	,477	,058	1,527	,165
	Huynh-Feldt	3,262	2,000	1,631	,763	,477	,058	1,527	,165
	Lower-bound	3,262	2,000	1,631	,763	,477	,058	1,527	,165
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	2,047	2	1,024	,479	,625	,037	,958	,120
	Greenhouse-Geisser	2,047	2,000	1,024	,479	,625	,037	,958	,120
	Huynh-Feldt	2,047	2,000	1,024	,479	,625	,037	,958	,120
	Lower-bound	2,047	2,000	1,024	,479	,625	,037	,958	,120
Error (Smile*Gaze*Proxemics)	Sphericity Assumed	53,421	25	2,137					
	Greenhouse-Geisser	53,421	25,000	2,137					
	Huynh-Feldt	53,421	25,000	2,137					
	Lower-bound	53,421	25,000	2,137					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentExtraversion

		95% Confidence Interval	
Mean	Std. Error	Lower Bound	Upper Bound
5,241	,150	4,933	5,550

Smile

Estimates

Measure: AgentExtraversion

			95% Confidence Interval	
Smile	Mean	Std. Error	Lower Bound	Upper Bound
1	5,143	,172	4,788	5,497
2	5,340	,190	4,949	5,731

Pairwise Comparisons

Measure: AgentExtraversion

					95% Confidence Interval for Difference[a]	
(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[a]	Lower Bound	Upper Bound
1	2	-,197	,204	,342	-,617	,222
2	1	,197	,204	,342	-,222	,617

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,036	,937[a]	1,000	25,000	,342	,036	,937	,154
Wilks' lambda	,964	,937[a]	1,000	25,000	,342	,036	,937	,154
Hotelling's trace	,037	,937[a]	1,000	25,000	,342	,036	,937	,154
Roy's largest root	,037	,937[a]	1,000	25,000	,342	,036	,937	,154

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze

Estimates

Measure: AgentExtraversion

Gaze	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,093	,176	4,730	5,455
2	5,390	,166	5,049	5,731

Pairwise Comparisons

Measure: AgentExtraversion

(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-,297	,164	,082	-,635	,041
2	1	,297	,164	,082	-,041	,635

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,116	3,281[a]	1,000	25,000	,082	,116	3,281	,414
Wilks' lambda	,884	3,281[a]	1,000	25,000	,082	,116	3,281	,414
Hotelling's trace	,131	3,281[a]	1,000	25,000	,082	,116	3,281	,414
Roy's largest root	,131	3,281[a]	1,000	25,000	,082	,116	3,281	,414

Each F tests the multivariate effect of Gaze. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Proxemics

Estimates

Measure: AgentExtraversion

			95% Confidence Interval	
Proxemics	Mean	Std. Error	Lower Bound	Upper Bound
1	4,600	,172	4,247	4,954
2	5,882	,191	5,490	6,275

Pairwise Comparisons

Measure: AgentExtraversion

					95% Confidence Interval for Difference^b	
(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.^{b}	Lower Bound	Upper Bound
1	2	-1,282*	,205	,000	-1,704	-,861
2	1	1,282*	,205	,000	,861	1,704

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Smile * Subject Agreeableness Factor

Estimates

Measure: AgentExtraversion

Smile	Subject Agreeableness Factor	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	4,758	,326	4,086	5,430
	MEDIUM	5,481	,315	4,833	6,129
	HIGH	5,189	,311	4,549	5,830
2	LOW	5,257	,360	4,515	5,999
	MEDIUM	4,933	,347	4,217	5,648
	HIGH	5,830	,343	5,123	6,537

Pairwise Comparisons

Measure: AgentExtraversion

Subject Agreeableness Factor	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
						Lower Bound	Upper Bound
LOW	1	2	-,499	,386	,208	-1,295	,297
	2	1	,499	,386	,208	-,297	1,295
MEDIUM	1	2	,548	,372	,154	-,219	1,315
	2	1	-,548	,372	,154	-1,315	,219
HIGH	1	2	-,641	,368	,094	-1,399	,117
	2	1	,641	,368	,094	-,117	1,399

Based on estimated marginal means

a. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Agreeableness Factor		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,063	1,667[a]	1,000	25,000	,208	,063	1,667	,237
	Wilks' lambda	,937	1,667[a]	1,000	25,000	,208	,063	1,667	,237
	Hotelling's trace	,067	1,667[a]	1,000	25,000	,208	,063	1,667	,237
	Roy's largest root	,067	1,667[a]	1,000	25,000	,208	,063	1,667	,237
MEDIUM	Pillai's trace	,080	2,166[a]	1,000	25,000	,154	,080	2,166	,293
	Wilks' lambda	,920	2,166[a]	1,000	25,000	,154	,080	2,166	,293
	Hotelling's trace	,087	2,166[a]	1,000	25,000	,154	,080	2,166	,293
	Roy's largest root	,087	2,166[a]	1,000	25,000	,154	,080	2,166	,293
HIGH	Pillai's trace	,108	3,031[a]	1,000	25,000	,094	,108	3,031	,388
	Wilks' lambda	,892	3,031[a]	1,000	25,000	,094	,108	3,031	,388
	Hotelling's trace	,121	3,031[a]	1,000	25,000	,094	,108	3,031	,388
	Roy's largest root	,121	3,031[a]	1,000	25,000	,094	,108	3,031	,388

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze * Subject Extraversion Factor

Estimates

Measure: AgentExtraversion

Gaze	Subject Extraversion Factor	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	4,922	,353	4,195	5,649
	MEDIUM	5,238	,331	4,556	5,920
	HIGH	5,119	,307	4,486	5,752
2	LOW	5,889	,333	5,204	6,574
	MEDIUM	5,327	,312	4,685	5,970
	HIGH	4,953	,290	4,357	5,550

Pairwise Comparisons

Measure: AgentExtraversion

Subject Extraversion Factor	(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
						Lower Bound	Upper Bound
LOW	1	2	-,968*	,329	,007	-1,646	-,289
	2	1	,968*	,329	,007	,289	1,646
MEDIUM	1	2	-,090	,309	,774	-,726	,547
	2	1	,090	,309	,774	-,547	,726
HIGH	1	2	,165	,287	,570	-,426	,756
	2	1	-,165	,287	,570	-,756	,426

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Extraversion Factor		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,257	8,626[a]	1,000	25,000	,007	,257	8,626	,806
	Wilks' lambda	,743	8,626[a]	1,000	25,000	,007	,257	8,626	,806
	Hotelling's trace	,345	8,626[a]	1,000	25,000	,007	,257	8,626	,806
	Roy's largest root	,345	8,626[a]	1,000	25,000	,007	,257	8,626	,806
MEDIUM	Pillai's trace	,003	,084[a]	1,000	25,000	,774	,003	,084	,059
	Wilks' lambda	,997	,084[a]	1,000	25,000	,774	,003	,084	,059
	Hotelling's trace	,003	,084[a]	1,000	25,000	,774	,003	,084	,059
	Roy's largest root	,003	,084[a]	1,000	25,000	,774	,003	,084	,059
HIGH	Pillai's trace	,013	,332[a]	1,000	25,000	,570	,013	,332	,086
	Wilks' lambda	,987	,332[a]	1,000	25,000	,570	,013	,332	,086
	Hotelling's trace	,013	,332[a]	1,000	25,000	,570	,013	,332	,086
	Roy's largest root	,013	,332[a]	1,000	25,000	,570	,013	,332	,086

Each F tests the multivariate simple effects of Gaze within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Agent Friendliness Measure

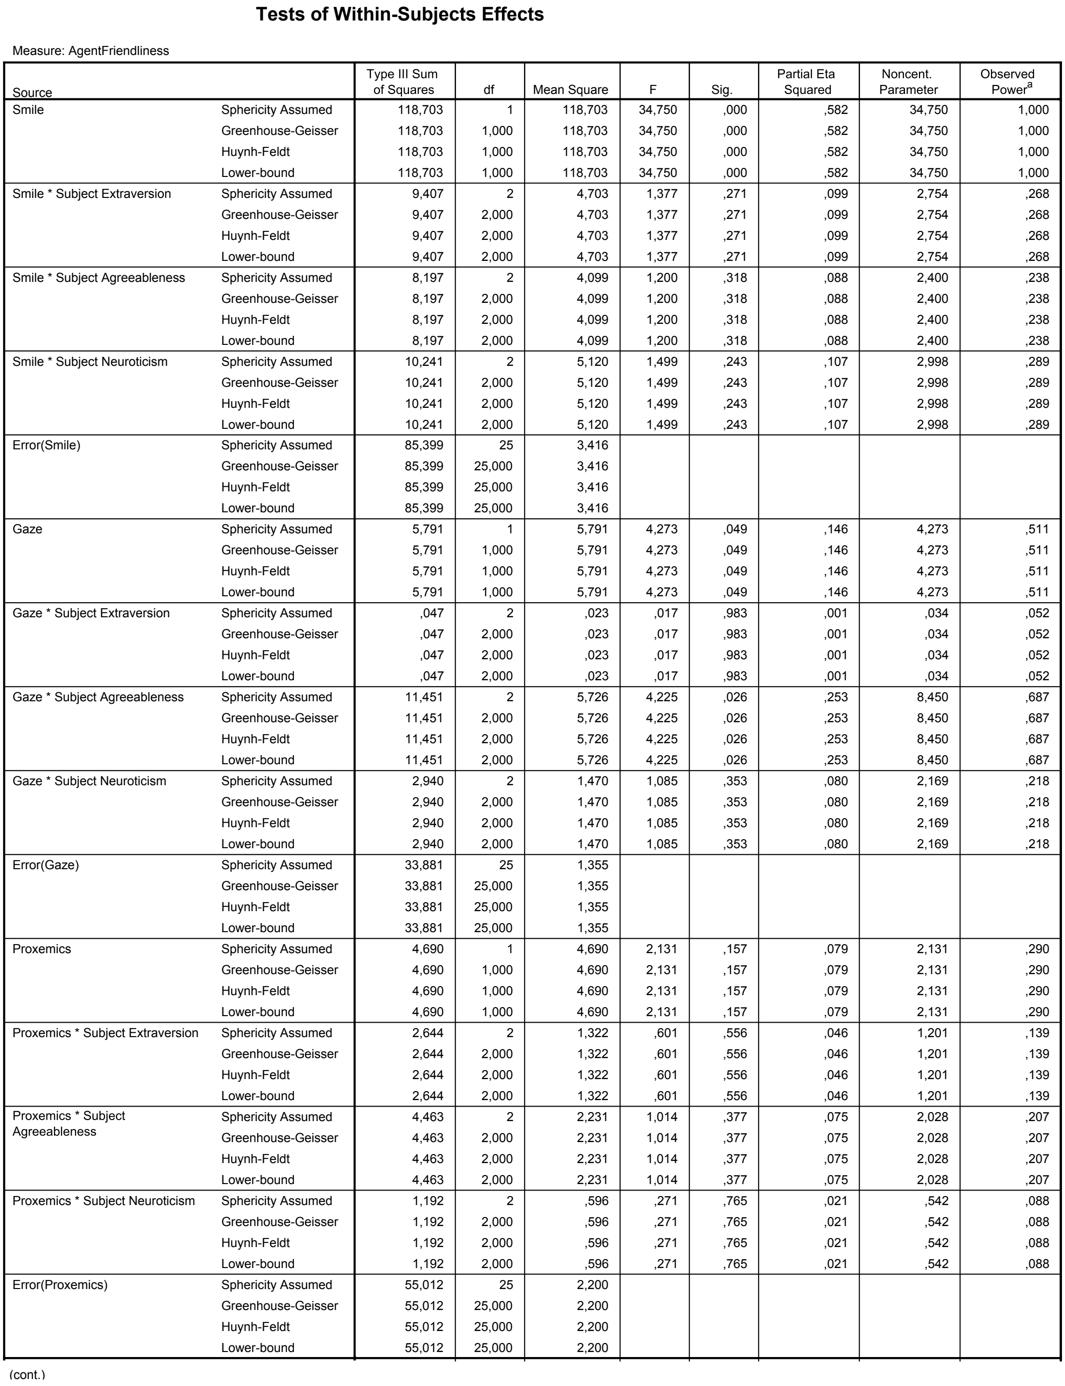

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power^a
Smile	Sphericity Assumed	118,703	1	118,703	34,750	,000	,582	34,750	1,000
	Greenhouse-Geisser	118,703	1,000	118,703	34,750	,000	,582	34,750	1,000
	Huynh-Feldt	118,703	1,000	118,703	34,750	,000	,582	34,750	1,000
	Lower-bound	118,703	1,000	118,703	34,750	,000	,582	34,750	1,000
Smile * Subject Extraversion	Sphericity Assumed	9,407	2	4,703	1,377	,271	,099	2,754	,268
	Greenhouse-Geisser	9,407	2,000	4,703	1,377	,271	,099	2,754	,268
	Huynh-Feldt	9,407	2,000	4,703	1,377	,271	,099	2,754	,268
	Lower-bound	9,407	2,000	4,703	1,377	,271	,099	2,754	,268
Smile * Subject Agreeableness	Sphericity Assumed	8,197	2	4,099	1,200	,318	,088	2,400	,238
	Greenhouse-Geisser	8,197	2,000	4,099	1,200	,318	,088	2,400	,238
	Huynh-Feldt	8,197	2,000	4,099	1,200	,318	,088	2,400	,238
	Lower-bound	8,197	2,000	4,099	1,200	,318	,088	2,400	,238
Smile * Subject Neuroticism	Sphericity Assumed	10,241	2	5,120	1,499	,243	,107	2,998	,289
	Greenhouse-Geisser	10,241	2,000	5,120	1,499	,243	,107	2,998	,289
	Huynh-Feldt	10,241	2,000	5,120	1,499	,243	,107	2,998	,289
	Lower-bound	10,241	2,000	5,120	1,499	,243	,107	2,998	,289
Error(Smile)	Sphericity Assumed	85,399	25	3,416					
	Greenhouse-Geisser	85,399	25,000	3,416					
	Huynh-Feldt	85,399	25,000	3,416					
	Lower-bound	85,399	25,000	3,416					
Gaze	Sphericity Assumed	5,791	1	5,791	4,273	,049	,146	4,273	,511
	Greenhouse-Geisser	5,791	1,000	5,791	4,273	,049	,146	4,273	,511
	Huynh-Feldt	5,791	1,000	5,791	4,273	,049	,146	4,273	,511
	Lower-bound	5,791	1,000	5,791	4,273	,049	,146	4,273	,511
Gaze * Subject Extraversion	Sphericity Assumed	,047	2	,023	,017	,983	,001	,034	,052
	Greenhouse-Geisser	,047	2,000	,023	,017	,983	,001	,034	,052
	Huynh-Feldt	,047	2,000	,023	,017	,983	,001	,034	,052
	Lower-bound	,047	2,000	,023	,017	,983	,001	,034	,052
Gaze * Subject Agreeableness	Sphericity Assumed	11,451	2	5,726	4,225	,026	,253	8,450	,687
	Greenhouse-Geisser	11,451	2,000	5,726	4,225	,026	,253	8,450	,687
	Huynh-Feldt	11,451	2,000	5,726	4,225	,026	,253	8,450	,687
	Lower-bound	11,451	2,000	5,726	4,225	,026	,253	8,450	,687
Gaze * Subject Neuroticism	Sphericity Assumed	2,940	2	1,470	1,085	,353	,080	2,169	,218
	Greenhouse-Geisser	2,940	2,000	1,470	1,085	,353	,080	2,169	,218
	Huynh-Feldt	2,940	2,000	1,470	1,085	,353	,080	2,169	,218
	Lower-bound	2,940	2,000	1,470	1,085	,353	,080	2,169	,218
Error(Gaze)	Sphericity Assumed	33,881	25	1,355					
	Greenhouse-Geisser	33,881	25,000	1,355					
	Huynh-Feldt	33,881	25,000	1,355					
	Lower-bound	33,881	25,000	1,355					
Proxemics	Sphericity Assumed	4,690	1	4,690	2,131	,157	,079	2,131	,290
	Greenhouse-Geisser	4,690	1,000	4,690	2,131	,157	,079	2,131	,290
	Huynh-Feldt	4,690	1,000	4,690	2,131	,157	,079	2,131	,290
	Lower-bound	4,690	1,000	4,690	2,131	,157	,079	2,131	,290
Proxemics * Subject Extraversion	Sphericity Assumed	2,644	2	1,322	,601	,556	,046	1,201	,139
	Greenhouse-Geisser	2,644	2,000	1,322	,601	,556	,046	1,201	,139
	Huynh-Feldt	2,644	2,000	1,322	,601	,556	,046	1,201	,139
	Lower-bound	2,644	2,000	1,322	,601	,556	,046	1,201	,139
Proxemics * Subject Agreeableness	Sphericity Assumed	4,463	2	2,231	1,014	,377	,075	2,028	,207
	Greenhouse-Geisser	4,463	2,000	2,231	1,014	,377	,075	2,028	,207
	Huynh-Feldt	4,463	2,000	2,231	1,014	,377	,075	2,028	,207
	Lower-bound	4,463	2,000	2,231	1,014	,377	,075	2,028	,207
Proxemics * Subject Neuroticism	Sphericity Assumed	1,192	2	,596	,271	,765	,021	,542	,088
	Greenhouse-Geisser	1,192	2,000	,596	,271	,765	,021	,542	,088
	Huynh-Feldt	1,192	2,000	,596	,271	,765	,021	,542	,088
	Lower-bound	1,192	2,000	,596	,271	,765	,021	,542	,088
Error(Proxemics)	Sphericity Assumed	55,012	25	2,200					
	Greenhouse-Geisser	55,012	25,000	2,200					
	Huynh-Feldt	55,012	25,000	2,200					
	Lower-bound	55,012	25,000	2,200					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power^a
Smile * Gaze	Sphericity Assumed	,334	1	,334	,168	,685	,007	,168	,068
	Greenhouse-Geisser	,334	1,000	,334	,168	,685	,007	,168	,068
	Huynh-Feldt	,334	1,000	,334	,168	,685	,007	,168	,068
	Lower-bound	,334	1,000	,334	,168	,685	,007	,168	,068
Smile * Gaze * Subject Extraversion	Sphericity Assumed	2,598	2	1,299	,656	,528	,050	1,312	,148
	Greenhouse-Geisser	2,598	2,000	1,299	,656	,528	,050	1,312	,148
	Huynh-Feldt	2,598	2,000	1,299	,656	,528	,050	1,312	,148
	Lower-bound	2,598	2,000	1,299	,656	,528	,050	1,312	,148
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	2,835	2	1,418	,716	,499	,054	1,431	,157
	Greenhouse-Geisser	2,835	2,000	1,418	,716	,499	,054	1,431	,157
	Huynh-Feldt	2,835	2,000	1,418	,716	,499	,054	1,431	,157
	Lower-bound	2,835	2,000	1,418	,716	,499	,054	1,431	,157
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	,390	2	,195	,098	,907	,008	,197	,063
	Greenhouse-Geisser	,390	2,000	,195	,098	,907	,008	,197	,063
	Huynh-Feldt	,390	2,000	,195	,098	,907	,008	,197	,063
	Lower-bound	,390	2,000	,195	,098	,907	,008	,197	,063
Error(Smile*Gaze)	Sphericity Assumed	49,512	25	1,980					
	Greenhouse-Geisser	49,512	25,000	1,980					
	Huynh-Feldt	49,512	25,000	1,980					
	Lower-bound	49,512	25,000	1,980					
Smile * Proxemics	Sphericity Assumed	,224	1	,224	,383	,542	,015	,383	,092
	Greenhouse-Geisser	,224	1,000	,224	,383	,542	,015	,383	,092
	Huynh-Feldt	,224	1,000	,224	,383	,542	,015	,383	,092
	Lower-bound	,224	1,000	,224	,383	,542	,015	,383	,092
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	3,128	2	1,564	2,681	,088	,177	5,362	,483
	Greenhouse-Geisser	3,128	2,000	1,564	2,681	,088	,177	5,362	,483
	Huynh-Feldt	3,128	2,000	1,564	2,681	,088	,177	5,362	,483
	Lower-bound	3,128	2,000	1,564	2,681	,088	,177	5,362	,483
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	4,695	2	2,347	4,023	,031	,243	8,046	,664
	Greenhouse-Geisser	4,695	2,000	2,347	4,023	,031	,243	8,046	,664
	Huynh-Feldt	4,695	2,000	2,347	4,023	,031	,243	8,046	,664
	Lower-bound	4,695	2,000	2,347	4,023	,031	,243	8,046	,664
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	,008	2	,004	,007	,993	,001	,015	,051
	Greenhouse-Geisser	,008	2,000	,004	,007	,993	,001	,015	,051
	Huynh-Feldt	,008	2,000	,004	,007	,993	,001	,015	,051
	Lower-bound	,008	2,000	,004	,007	,993	,001	,015	,051
Error(Smile*Proxemics)	Sphericity Assumed	14,587	25	,583					
	Greenhouse-Geisser	14,587	25,000	,583					
	Huynh-Feldt	14,587	25,000	,583					
	Lower-bound	14,587	25,000	,583					
Gaze * Proxemics	Sphericity Assumed	,020	1	,020	,025	,877	,001	,025	,053
	Greenhouse-Geisser	,020	1,000	,020	,025	,877	,001	,025	,053
	Huynh-Feldt	,020	1,000	,020	,025	,877	,001	,025	,053
	Lower-bound	,020	1,000	,020	,025	,877	,001	,025	,053
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	2,192	2	1,096	1,319	,285	,095	2,638	,258
	Greenhouse-Geisser	2,192	2,000	1,096	1,319	,285	,095	2,638	,258
	Huynh-Feldt	2,192	2,000	1,096	1,319	,285	,095	2,638	,258
	Lower-bound	2,192	2,000	1,096	1,319	,285	,095	2,638	,258
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	1,118	2	,559	,673	,519	,051	1,346	,150
	Greenhouse-Geisser	1,118	2,000	,559	,673	,519	,051	1,346	,150
	Huynh-Feldt	1,118	2,000	,559	,673	,519	,051	1,346	,150
	Lower-bound	1,118	2,000	,559	,673	,519	,051	1,346	,150
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,131	2	,066	,079	,924	,006	,158	,061
	Greenhouse-Geisser	,131	2,000	,066	,079	,924	,006	,158	,061
	Huynh-Feldt	,131	2,000	,066	,079	,924	,006	,158	,061
	Lower-bound	,131	2,000	,066	,079	,924	,006	,158	,061
Error(Gaze*Proxemics)	Sphericity Assumed	20,777	25	,831					
	Greenhouse-Geisser	20,777	25,000	,831					
	Huynh-Feldt	20,777	25,000	,831					
	Lower-bound	20,777	25,000	,831					
Smile * Gaze * Proxemics	Sphericity Assumed	,319	1	,319	,321	,576	,013	,321	,085
	Greenhouse-Geisser	,319	1,000	,319	,321	,576	,013	,321	,085

(cont.)

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	,319	1,000	,319	,321	,576	,013	,321	,085
	Lower-bound	,319	1,000	,319	,321	,576	,013	,321	,085
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	2,913	2	1,456	1,466	,250	,105	2,932	,283
	Greenhouse-Geisser	2,913	2,000	1,456	1,466	,250	,105	2,932	,283
	Huynh-Feldt	2,913	2,000	1,456	1,466	,250	,105	2,932	,283
	Lower-bound	2,913	2,000	1,456	1,466	,250	,105	2,932	,283
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	2,812	2	1,406	1,416	,262	,102	2,831	,275
	Greenhouse-Geisser	2,812	2,000	1,406	1,416	,262	,102	2,831	,275
	Huynh-Feldt	2,812	2,000	1,406	1,416	,262	,102	2,831	,275
	Lower-bound	2,812	2,000	1,406	1,416	,262	,102	2,831	,275
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	5,201	2	2,601	2,618	,093	,173	5,236	,473
	Greenhouse-Geisser	5,201	2,000	2,601	2,618	,093	,173	5,236	,473
	Huynh-Feldt	5,201	2,000	2,601	2,618	,093	,173	5,236	,473
	Lower-bound	5,201	2,000	2,601	2,618	,093	,173	5,236	,473
Error(Smile*Gaze*Proxemics)	Sphericity Assumed	24,832	25	,993					
	Greenhouse-Geisser	24,832	25,000	,993					
	Huynh-Feldt	24,832	25,000	,993					
	Lower-bound	24,832	25,000	,993					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentFriendliness

		95% Confidence Interval	
Mean	Std. Error	Lower Bound	Upper Bound
5,816	,154	5,499	6,133

Smile

Estimates

Measure: AgentFriendliness

			95% Confidence Interval	
Smile	Mean	Std. Error	Lower Bound	Upper Bound
1	5,124	,164	4,786	5,461
2	6,508	,219	6,057	6,960

Pairwise Comparisons

Measure: AgentFriendliness

					95% Confidence Interval for Difference[b]	
(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	Lower Bound	Upper Bound
1	2	-1,384[*]	,235	,000	-1,868	-,901
2	1	1,384[*]	,235	,000	,901	1,868

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,582	34,750[a]	1,000	25,000	,000	,582	34,750	1,000
Wilks' lambda	,418	34,750[a]	1,000	25,000	,000	,582	34,750	1,000
Hotelling's trace	1,390	34,750[a]	1,000	25,000	,000	,582	34,750	1,000
Roy's largest root	1,390	34,750[a]	1,000	25,000	,000	,582	34,750	1,000

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze

Estimates

Measure: AgentFriendliness

Gaze	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,663	,179	5,294	6,032
2	5,969	,162	5,635	6,302

Pairwise Comparisons

Measure: AgentFriendliness

(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
					Lower Bound	Upper Bound
1	2	-,306*	,148	,049	-,610	-,001
2	1	,306*	,148	,049	,001	,610

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,146	4,273[a]	1,000	25,000	,049	,146	4,273	,511
Wilks' lambda	,854	4,273[a]	1,000	25,000	,049	,146	4,273	,511
Hotelling's trace	,171	4,273[a]	1,000	25,000	,049	,146	4,273	,511
Roy's largest root	,171	4,273[a]	1,000	25,000	,049	,146	4,273	,511

Each F tests the multivariate effect of Gaze. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Proxemics

Estimates

Measure: AgentFriendliness

Proxemics	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,678	,172	5,325	6,032
2	5,954	,189	5,565	6,343

Pairwise Comparisons

Measure: AgentFriendliness

(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-,275	,188	,157	-,663	,113
2	1	,275	,188	,157	-,113	,663

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,079	2,131[a]	1,000	25,000	,157	,079	2,131	,290
Wilks' lambda	,921	2,131[a]	1,000	25,000	,157	,079	2,131	,290
Hotelling's trace	,085	2,131[a]	1,000	25,000	,157	,079	2,131	,290
Roy's largest root	,085	2,131[a]	1,000	25,000	,157	,079	2,131	,290

Each F tests the multivariate effect of Proxemics. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Subject Agreeableness Factor * Gaze

Estimates

Measure: AgentFriendliness

Subject Agreeableness Factor	Gaze	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
LOW	1	5,056	,340	4,357	5,756
	2	5,963	,307	5,331	6,596
MEDIUM	1	6,245	,327	5,571	6,919
	2	6,469	,296	5,860	7,079
HIGH	1	5,688	,324	5,021	6,354
	2	5,474	,293	4,871	6,077

Pairwise Comparisons

Measure: AgentFriendliness

Subject Agreeableness Factor	(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
						Lower Bound	Upper Bound
LOW	1	2	-,907*	,281	,003	-1,485	-,329
	2	1	,907*	,281	,003	,329	1,485
MEDIUM	1	2	-,224	,270	,415	-,781	,333
	2	1	,224	,270	,415	-,333	,781
HIGH	1	2	,214	,267	,432	-,337	,764
	2	1	-,214	,267	,432	-,764	,337

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Agreeableness Factor		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,295	10,446[a]	1,000	25,000	,003	,295	10,446	,874
	Wilks' lambda	,705	10,446[a]	1,000	25,000	,003	,295	10,446	,874
	Hotelling's trace	,418	10,446[a]	1,000	25,000	,003	,295	10,446	,874
	Roy's largest root	,418	10,446[a]	1,000	25,000	,003	,295	10,446	,874
MEDIUM	Pillai's trace	,027	,686[a]	1,000	25,000	,415	,027	,686	,125
	Wilks' lambda	,973	,686[a]	1,000	25,000	,415	,027	,686	,125
	Hotelling's trace	,027	,686[a]	1,000	25,000	,415	,027	,686	,125
	Roy's largest root	,027	,686[a]	1,000	25,000	,415	,027	,686	,125
HIGH	Pillai's trace	,025	,638[a]	1,000	25,000	,432	,025	,638	,120
	Wilks' lambda	,975	,638[a]	1,000	25,000	,432	,025	,638	,120
	Hotelling's trace	,026	,638[a]	1,000	25,000	,432	,025	,638	,120
	Roy's largest root	,026	,638[a]	1,000	25,000	,432	,025	,638	,120

Each F tests the multivariate simple effects of Gaze within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Subject Agreeableness Factor * Proxemics * Smile

Estimates

Measure: AgentFriendliness

Subject Agreeableness Factor	Proxemics	Smile	Mean	Std. Error	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
LOW	1	1	5,238	,421	4,372	6,105
		2	5,707	,405	4,874	6,541
	2	1	4,844	,321	4,183	5,504
		2	6,250	,507	5,205	7,295
MEDIUM	1	1	5,197	,405	4,362	6,032
		2	7,437	,390	6,633	8,240
	2	1	5,577	,309	4,940	6,213
		2	7,219	,489	6,212	8,226
HIGH	1	1	4,614	,401	3,788	5,439
		2	5,878	,386	5,084	6,672
	2	1	5,274	,305	4,645	5,903
		2	6,558	,483	5,563	7,554

Pairwise Comparisons

Measure: AgentFriendliness

Subject Agreeableness Factor	Proxemics	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b] Lower Bound	95% Confidence Interval for Difference[b] Upper Bound
LOW	1	1	2	-,469	,507	,364	-1,514	,576
		2	1	,469	,507	,364	-,576	1,514
	2	1	2	-1,407*	,455	,005	-2,344	-,469
		2	1	1,407*	,455	,005	,469	2,344
MEDIUM	1	1	2	-2,240*	,489	,000	-3,247	-1,232
		2	1	2,240*	,489	,000	1,232	3,247
	2	1	2	-1,642*	,439	,001	-2,546	-,739
		2	1	1,642*	,439	,001	,739	2,546
HIGH	1	1	2	-1,264*	,483	,015	-2,260	-,268
		2	1	1,264*	,483	,015	,268	2,260
	2	1	2	-1,284*	,434	,007	-2,177	-,391
		2	1	1,284*	,434	,007	,391	2,177

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Agreeableness Factor	Proxemics		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	1	Pillai's trace	,033	,855[a]	1,000	25,000	,364	,033	,855	,144
		Wilks' lambda	,967	,855[a]	1,000	25,000	,364	,033	,855	,144
		Hotelling's trace	,034	,855[a]	1,000	25,000	,364	,033	,855	,144
		Roy's largest root	,034	,855[a]	1,000	25,000	,364	,033	,855	,144
	2	Pillai's trace	,276	9,552[a]	1,000	25,000	,005	,276	9,552	,844
		Wilks' lambda	,724	9,552[a]	1,000	25,000	,005	,276	9,552	,844
		Hotelling's trace	,382	9,552[a]	1,000	25,000	,005	,276	9,552	,844
		Roy's largest root	,382	9,552[a]	1,000	25,000	,005	,276	9,552	,844
MEDIUM	1	Pillai's trace	,456	20,971[a]	1,000	25,000	,000	,456	20,971	,993
		Wilks' lambda	,544	20,971[a]	1,000	25,000	,000	,456	20,971	,993
		Hotelling's trace	,839	20,971[a]	1,000	25,000	,000	,456	20,971	,993
		Roy's largest root	,839	20,971[a]	1,000	25,000	,000	,456	20,971	,993
	2	Pillai's trace	,359	14,015[a]	1,000	25,000	,001	,359	14,015	,949
		Wilks' lambda	,641	14,015[a]	1,000	25,000	,001	,359	14,015	,949
		Hotelling's trace	,561	14,015[a]	1,000	25,000	,001	,359	14,015	,949
		Roy's largest root	,561	14,015[a]	1,000	25,000	,001	,359	14,015	,949
HIGH	1	Pillai's trace	,215	6,837[a]	1,000	25,000	,015	,215	6,837	,710
		Wilks' lambda	,785	6,837[a]	1,000	25,000	,015	,215	6,837	,710
		Hotelling's trace	,273	6,837[a]	1,000	25,000	,015	,215	6,837	,710
		Roy's largest root	,273	6,837[a]	1,000	25,000	,015	,215	6,837	,710
	2	Pillai's trace	,260	8,772[a]	1,000	25,000	,007	,260	8,772	,812
		Wilks' lambda	,740	8,772[a]	1,000	25,000	,007	,260	8,772	,812
		Hotelling's trace	,351	8,772[a]	1,000	25,000	,007	,260	8,772	,812
		Roy's largest root	,351	8,772[a]	1,000	25,000	,007	,260	8,772	,812

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Subject Agreeableness Factor * Smile * Proxemics

Estimates

Measure: AgentFriendliness

Subject Agreeableness Factor	Smile	Proxemics	Mean	Std. Error	95% Confidence Interval	
					Lower Bound	Upper Bound
LOW	1	1	5,238	,421	4,372	6,105
		2	4,844	,321	4,183	5,504
	2	1	5,707	,405	4,874	6,541
		2	6,250	,507	5,205	7,295
MEDIUM	1	1	5,197	,405	4,362	6,032
		2	5,577	,309	4,940	6,213
	2	1	7,437	,390	6,633	8,240
		2	7,219	,489	6,212	8,226
HIGH	1	1	4,614	,401	3,788	5,439
		2	5,274	,305	4,645	5,903
	2	1	5,878	,386	5,084	6,672
		2	6,558	,483	5,563	7,554

Pairwise Comparisons

Measure: AgentFriendliness

Subject Agreeableness Factor	Smile	(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence [a]...
							Lower Bound
LOW	1	1	2	,395	,416	,352	-,463
		2	1	-,395	,416	,352	-1,252
	2	1	2	-,543	,388	,173	-1,341
		2	1	,543	,388	,173	-,255
MEDIUM	1	1	2	-,380	,401	,353	-1,206
		2	1	,380	,401	,353	-,447
	2	1	2	,218	,373	,565	-,552
		2	1	-,218	,373	,565	-,987
HIGH	1	1	2	-,660	,397	,108	-1,477
		2	1	,660	,397	,108	-,157
	2	1	2	-,680	,369	,077	-1,441
		2	1	,680	,369	,077	-,080

Pairwise Comparisons

Measure: AgentFriendliness

Subject Agreeableness Factor	Smile	(I) Proxemics	(J) Proxemics	95% Confidence [a]... Upper Bound
LOW	1	1	2	1,252
		2	1	,463
	2	1	2	,255
		2	1	1,341
MEDIUM	1	1	2	,447
		2	1	1,206
	2	1	2	,987
		2	1	,552
HIGH	1	1	2	,157
		2	1	1,477
	2	1	2	,080
		2	1	1,441

Based on estimated marginal means

a. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Agreeableness Factor	Smile		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	1	Pillai's trace	,035	,899[a]	1,000	25,000	,352	,035	,899	,149
		Wilks' lambda	,965	,899[a]	1,000	25,000	,352	,035	,899	,149
		Hotelling's trace	,036	,899[a]	1,000	25,000	,352	,035	,899	,149
		Roy's largest root	,036	,899[a]	1,000	25,000	,352	,035	,899	,149
	2	Pillai's trace	,073	1,963[a]	1,000	25,000	,173	,073	1,963	,271
		Wilks' lambda	,927	1,963[a]	1,000	25,000	,173	,073	1,963	,271
		Hotelling's trace	,079	1,963[a]	1,000	25,000	,173	,073	1,963	,271
		Roy's largest root	,079	1,963[a]	1,000	25,000	,173	,073	1,963	,271
MEDIUM	1	Pillai's trace	,035	,896[a]	1,000	25,000	,353	,035	,896	,149
		Wilks' lambda	,965	,896[a]	1,000	25,000	,353	,035	,896	,149
		Hotelling's trace	,036	,896[a]	1,000	25,000	,353	,035	,896	,149
		Roy's largest root	,036	,896[a]	1,000	25,000	,353	,035	,896	,149
	2	Pillai's trace	,013	,340[a]	1,000	25,000	,565	,013	,340	,087
		Wilks' lambda	,987	,340[a]	1,000	25,000	,565	,013	,340	,087
		Hotelling's trace	,014	,340[a]	1,000	25,000	,565	,013	,340	,087
		Roy's largest root	,014	,340[a]	1,000	25,000	,565	,013	,340	,087
HIGH	1	Pillai's trace	,100	2,771[a]	1,000	25,000	,108	,100	2,771	,360
		Wilks' lambda	,900	2,771[a]	1,000	25,000	,108	,100	2,771	,360
		Hotelling's trace	,111	2,771[a]	1,000	25,000	,108	,100	2,771	,360
		Roy's largest root	,111	2,771[a]	1,000	25,000	,108	,100	2,771	,360
	2	Pillai's trace	,120	3,397[a]	1,000	25,000	,077	,120	3,397	,426
		Wilks' lambda	,880	3,397[a]	1,000	25,000	,077	,120	3,397	,426
		Hotelling's trace	,136	3,397[a]	1,000	25,000	,077	,120	3,397	,426
		Roy's largest root	,136	3,397[a]	1,000	25,000	,077	,120	3,397	,426

Each F tests the multivariate simple effects of Proxemics within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Agent Likeability Measure

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Smile	Sphericity Assumed	22,124	1	22,124	20,034	,000	,445	20,034	,990
	Greenhouse-Geisser	22,124	1,000	22,124	20,034	,000	,445	20,034	,990
	Huynh-Feldt	22,124	1,000	22,124	20,034	,000	,445	20,034	,990
	Lower-bound	22,124	1,000	22,124	20,034	,000	,445	20,034	,990
Smile * Subject Extraversion	Sphericity Assumed	,615	2	,308	,279	,759	,022	,557	,089
	Greenhouse-Geisser	,615	2,000	,308	,279	,759	,022	,557	,089
	Huynh-Feldt	,615	2,000	,308	,279	,759	,022	,557	,089
	Lower-bound	,615	2,000	,308	,279	,759	,022	,557	,089
Smile * Subject Agreeableness	Sphericity Assumed	4,217	2	2,108	1,909	,169	,133	3,819	,358
	Greenhouse-Geisser	4,217	2,000	2,108	1,909	,169	,133	3,819	,358
	Huynh-Feldt	4,217	2,000	2,108	1,909	,169	,133	3,819	,358
	Lower-bound	4,217	2,000	2,108	1,909	,169	,133	3,819	,358
Smile * Subject Neuroticism	Sphericity Assumed	4,637	2	2,319	2,100	,144	,144	4,199	,390
	Greenhouse-Geisser	4,637	2,000	2,319	2,100	,144	,144	4,199	,390
	Huynh-Feldt	4,637	2,000	2,319	2,100	,144	,144	4,199	,390
	Lower-bound	4,637	2,000	2,319	2,100	,144	,144	4,199	,390
Error(Smile)	Sphericity Assumed	27,608	25	1,104					
	Greenhouse-Geisser	27,608	25,000	1,104					
	Huynh-Feldt	27,608	25,000	1,104					
	Lower-bound	27,608	25,000	1,104					
Gaze	Sphericity Assumed	1,148	1	1,148	1,832	,188	,068	1,832	,256
	Greenhouse-Geisser	1,148	1,000	1,148	1,832	,188	,068	1,832	,256
	Huynh-Feldt	1,148	1,000	1,148	1,832	,188	,068	1,832	,256
	Lower-bound	1,148	1,000	1,148	1,832	,188	,068	1,832	,256
Gaze * Subject Extraversion	Sphericity Assumed	,048	2	,024	,038	,963	,003	,076	,055
	Greenhouse-Geisser	,048	2,000	,024	,038	,963	,003	,076	,055
	Huynh-Feldt	,048	2,000	,024	,038	,963	,003	,076	,055
	Lower-bound	,048	2,000	,024	,038	,963	,003	,076	,055
Gaze * Subject Agreeableness	Sphericity Assumed	2,919	2	1,459	2,328	,118	,157	4,656	,427
	Greenhouse-Geisser	2,919	2,000	1,459	2,328	,118	,157	4,656	,427
	Huynh-Feldt	2,919	2,000	1,459	2,328	,118	,157	4,656	,427
	Lower-bound	2,919	2,000	1,459	2,328	,118	,157	4,656	,427
Gaze * Subject Neuroticism	Sphericity Assumed	,504	2	,252	,402	,673	,031	,805	,108
	Greenhouse-Geisser	,504	2,000	,252	,402	,673	,031	,805	,108
	Huynh-Feldt	,504	2,000	,252	,402	,673	,031	,805	,108
	Lower-bound	,504	2,000	,252	,402	,673	,031	,805	,108
Error(Gaze)	Sphericity Assumed	15,670	25	,627					
	Greenhouse-Geisser	15,670	25,000	,627					
	Huynh-Feldt	15,670	25,000	,627					
	Lower-bound	15,670	25,000	,627					
Proxemics	Sphericity Assumed	1,291	1	1,291	1,389	,250	,053	1,389	,205
	Greenhouse-Geisser	1,291	1,000	1,291	1,389	,250	,053	1,389	,205
	Huynh-Feldt	1,291	1,000	1,291	1,389	,250	,053	1,389	,205
	Lower-bound	1,291	1,000	1,291	1,389	,250	,053	1,389	,205
Proxemics * Subject Extraversion	Sphericity Assumed	1,861	2	,931	1,001	,382	,074	2,002	,204
	Greenhouse-Geisser	1,861	2,000	,931	1,001	,382	,074	2,002	,204
	Huynh-Feldt	1,861	2,000	,931	1,001	,382	,074	2,002	,204
	Lower-bound	1,861	2,000	,931	1,001	,382	,074	2,002	,204
Proxemics * Subject Agreeableness	Sphericity Assumed	1,827	2	,914	,983	,388	,073	1,966	,201
	Greenhouse-Geisser	1,827	2,000	,914	,983	,388	,073	1,966	,201
	Huynh-Feldt	1,827	2,000	,914	,983	,388	,073	1,966	,201
	Lower-bound	1,827	2,000	,914	,983	,388	,073	1,966	,201
Proxemics * Subject Neuroticism	Sphericity Assumed	1,649	2	,825	,887	,424	,066	1,774	,185
	Greenhouse-Geisser	1,649	2,000	,825	,887	,424	,066	1,774	,185
	Huynh-Feldt	1,649	2,000	,825	,887	,424	,066	1,774	,185
	Lower-bound	1,649	2,000	,825	,887	,424	,066	1,774	,185
Error(Proxemics)	Sphericity Assumed	23,237	25	,929					
	Greenhouse-Geisser	23,237	25,000	,929					
	Huynh-Feldt	23,237	25,000	,929					
	Lower-bound	23,237	25,000	,929					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Smile * Gaze	Sphericity Assumed	,242	1	,242	,425	,520	,017	,425	,096
	Greenhouse-Geisser	,242	1,000	,242	,425	,520	,017	,425	,096
	Huynh-Feldt	,242	1,000	,242	,425	,520	,017	,425	,096
	Lower-bound	,242	1,000	,242	,425	,520	,017	,425	,096
Smile * Gaze * Subject Extraversion	Sphericity Assumed	,784	2	,392	,689	,511	,052	1,378	,153
	Greenhouse-Geisser	,784	2,000	,392	,689	,511	,052	1,378	,153
	Huynh-Feldt	,784	2,000	,392	,689	,511	,052	1,378	,153
	Lower-bound	,784	2,000	,392	,689	,511	,052	1,378	,153
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	2,156	2	1,078	1,894	,171	,132	3,789	,356
	Greenhouse-Geisser	2,156	2,000	1,078	1,894	,171	,132	3,789	,356
	Huynh-Feldt	2,156	2,000	1,078	1,894	,171	,132	3,789	,356
	Lower-bound	2,156	2,000	1,078	1,894	,171	,132	3,789	,356
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	,085	2	,042	,075	,928	,006	,149	,060
	Greenhouse-Geisser	,085	2,000	,042	,075	,928	,006	,149	,060
	Huynh-Feldt	,085	2,000	,042	,075	,928	,006	,149	,060
	Lower-bound	,085	2,000	,042	,075	,928	,006	,149	,060
Error(Smile*Gaze)	Sphericity Assumed	14,226	25	,569					
	Greenhouse-Geisser	14,226	25,000	,569					
	Huynh-Feldt	14,226	25,000	,569					
	Lower-bound	14,226	25,000	,569					
Smile * Proxemics	Sphericity Assumed	,300	1	,300	,385	,540	,015	,385	,092
	Greenhouse-Geisser	,300	1,000	,300	,385	,540	,015	,385	,092
	Huynh-Feldt	,300	1,000	,300	,385	,540	,015	,385	,092
	Lower-bound	,300	1,000	,300	,385	,540	,015	,385	,092
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	,716	2	,358	,460	,636	,036	,920	,117
	Greenhouse-Geisser	,716	2,000	,358	,460	,636	,036	,920	,117
	Huynh-Feldt	,716	2,000	,358	,460	,636	,036	,920	,117
	Lower-bound	,716	2,000	,358	,460	,636	,036	,920	,117
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	1,781	2	,890	1,145	,334	,084	2,290	,229
	Greenhouse-Geisser	1,781	2,000	,890	1,145	,334	,084	2,290	,229
	Huynh-Feldt	1,781	2,000	,890	1,145	,334	,084	2,290	,229
	Lower-bound	1,781	2,000	,890	1,145	,334	,084	2,290	,229
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	1,020	2	,510	,656	,528	,050	1,312	,148
	Greenhouse-Geisser	1,020	2,000	,510	,656	,528	,050	1,312	,148
	Huynh-Feldt	1,020	2,000	,510	,656	,528	,050	1,312	,148
	Lower-bound	1,020	2,000	,510	,656	,528	,050	1,312	,148
Error(Smile*Proxemics)	Sphericity Assumed	19,438	25	,778					
	Greenhouse-Geisser	19,438	25,000	,778					
	Huynh-Feldt	19,438	25,000	,778					
	Lower-bound	19,438	25,000	,778					
Gaze * Proxemics	Sphericity Assumed	,203	1	,203	,626	,436	,024	,626	,119
	Greenhouse-Geisser	,203	1,000	,203	,626	,436	,024	,626	,119
	Huynh-Feldt	,203	1,000	,203	,626	,436	,024	,626	,119
	Lower-bound	,203	1,000	,203	,626	,436	,024	,626	,119
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	,891	2	,445	1,377	,271	,099	2,753	,268
	Greenhouse-Geisser	,891	2,000	,445	1,377	,271	,099	2,753	,268
	Huynh-Feldt	,891	2,000	,445	1,377	,271	,099	2,753	,268
	Lower-bound	,891	2,000	,445	1,377	,271	,099	2,753	,268
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	1,833	2	,916	2,832	,078	,185	5,664	,505
	Greenhouse-Geisser	1,833	2,000	,916	2,832	,078	,185	5,664	,505
	Huynh-Feldt	1,833	2,000	,916	2,832	,078	,185	5,664	,505
	Lower-bound	1,833	2,000	,916	2,832	,078	,185	5,664	,505
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,700	2	,350	1,081	,355	,080	2,163	,218
	Greenhouse-Geisser	,700	2,000	,350	1,081	,355	,080	2,163	,218
	Huynh-Feldt	,700	2,000	,350	1,081	,355	,080	2,163	,218
	Lower-bound	,700	2,000	,350	1,081	,355	,080	2,163	,218
Error(Gaze*Proxemics)	Sphericity Assumed	8,089	25	,324					
	Greenhouse-Geisser	8,089	25,000	,324					
	Huynh-Feldt	8,089	25,000	,324					
	Lower-bound	8,089	25,000	,324					
Smile * Gaze * Proxemics	Sphericity Assumed	,074	1	,074	,135	,716	,005	,135	,064
	Greenhouse-Geisser	,074	1,000	,074	,135	,716	,005	,135	,064

(cont.)

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	,074	1,000	,074	,135	,716	,005	,135	,064
	Lower-bound	,074	1,000	,074	,135	,716	,005	,135	,064
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	,190	2	,095	,173	,842	,014	,346	,074
	Greenhouse-Geisser	,190	2,000	,095	,173	,842	,014	,346	,074
	Huynh-Feldt	,190	2,000	,095	,173	,842	,014	,346	,074
	Lower-bound	,190	2,000	,095	,173	,842	,014	,346	,074
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	,151	2	,075	,137	,873	,011	,274	,069
	Greenhouse-Geisser	,151	2,000	,075	,137	,873	,011	,274	,069
	Huynh-Feldt	,151	2,000	,075	,137	,873	,011	,274	,069
	Lower-bound	,151	2,000	,075	,137	,873	,011	,274	,069
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,694	2	,347	,632	,540	,048	1,263	,144
	Greenhouse-Geisser	,694	2,000	,347	,632	,540	,048	1,263	,144
	Huynh-Feldt	,694	2,000	,347	,632	,540	,048	1,263	,144
	Lower-bound	,694	2,000	,347	,632	,540	,048	1,263	,144
Error(Smile*Gaze*Proxemics)	Sphericity Assumed	13,744	25	,550					
	Greenhouse-Geisser	13,744	25,000	,550					
	Huynh-Feldt	13,744	25,000	,550					
	Lower-bound	13,744	25,000	,550					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentLikeability

Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound
3,542	,100	3,336	3,748

Smile

Estimates

Measure: AgentLikeability

Smile	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	3,243	,106	3,024	3,463
2	3,841	,133	3,567	4,115

Pairwise Comparisons

Measure: AgentLikeability

(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
					Lower Bound	Upper Bound
1	2	-,598[*]	,134	,000	-,873	-,323
2	1	,598[*]	,134	,000	,323	,873

Based on estimated marginal means

[*]. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,445	20,034[a]	1,000	25,000	,000	,445	20,034	,990
Wilks' lambda	,555	20,034[a]	1,000	25,000	,000	,445	20,034	,990
Hotelling's trace	,801	20,034[a]	1,000	25,000	,000	,445	20,034	,990
Roy's largest root	,801	20,034[a]	1,000	25,000	,000	,445	20,034	,990

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze

Estimates

Measure: AgentLikeability

Gaze	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	3,474	,120	3,227	3,721
2	3,610	,104	3,397	3,824

Pairwise Comparisons

Measure: AgentLikeability

(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-,136	,101	,188	-,343	,071
2	1	,136	,101	,188	-,071	,343

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,068	1,832[a]	1,000	25,000	,188	,068	1,832	,256
Wilks' lambda	,932	1,832[a]	1,000	25,000	,188	,068	1,832	,256
Hotelling's trace	,073	1,832[a]	1,000	25,000	,188	,068	1,832	,256
Roy's largest root	,073	1,832[a]	1,000	25,000	,188	,068	1,832	,256

Each F tests the multivariate effect of Gaze. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Proxemics

Estimates

Measure: AgentLikeability

			95% Confidence Interval	
Proximity	Mean	Std. Error	Lower Bound	Upper Bound
1	3,470	,116	3,230	3,710
2	3,614	,118	3,371	3,858

Pairwise Comparisons

Measure: AgentLikeability

					95% Confidence Interval for Difference[a]	
(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[a]	Lower Bound	Upper Bound
1	2	-,144	,123	,250	-,397	,108
2	1	,144	,123	,250	-,108	,397

Based on estimated marginal means

a. Adjustment for multiple comparisons: Least Significant Difference (equivalent to no adjustments).

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,053	1,389[a]	1,000	25,000	,250	,053	1,389	,205
Wilks' lambda	,947	1,389[a]	1,000	25,000	,250	,053	1,389	,205
Hotelling's trace	,056	1,389[a]	1,000	25,000	,250	,053	1,389	,205
Roy's largest root	,056	1,389[a]	1,000	25,000	,250	,053	1,389	,205

Each F tests the multivariate effect of Proxemics. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Third Person Perspective Trial

Agent Extraversion Measure

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power^a
Smile	Sphericity Assumed	,268	1	,268	,383	,542	,015	,383	,091
	Greenhouse-Geisser	,268	1,000	,268	,383	,542	,015	,383	,091
	Huynh-Feldt	,268	1,000	,268	,383	,542	,015	,383	,091
	Lower-bound	,268	1,000	,268	,383	,542	,015	,383	,091
Smile * Subject Extraversion	Sphericity Assumed	5,982	2	2,991	4,272	,025	,255	8,545	,692
	Greenhouse-Geisser	5,982	2,000	2,991	4,272	,025	,255	8,545	,692
	Huynh-Feldt	5,982	2,000	2,991	4,272	,025	,255	8,545	,692
	Lower-bound	5,982	2,000	2,991	4,272	,025	,255	8,545	,692
Smile * Subject Agreeableness	Sphericity Assumed	2,882	2	1,441	2,058	,149	,141	4,117	,383
	Greenhouse-Geisser	2,882	2,000	1,441	2,058	,149	,141	4,117	,383
	Huynh-Feldt	2,882	2,000	1,441	2,058	,149	,141	4,117	,383
	Lower-bound	2,882	2,000	1,441	2,058	,149	,141	4,117	,383
Smile * Subject Neuroticism	Sphericity Assumed	1,516	2	,758	1,083	,354	,080	2,166	,218
	Greenhouse-Geisser	1,516	2,000	,758	1,083	,354	,080	2,166	,218
	Huynh-Feldt	1,516	2,000	,758	1,083	,354	,080	2,166	,218
	Lower-bound	1,516	2,000	,758	1,083	,354	,080	2,166	,218
Error(Smile)	Sphericity Assumed	17,501	25	,700					
	Greenhouse-Geisser	17,501	25,000	,700					
	Huynh-Feldt	17,501	25,000	,700					
	Lower-bound	17,501	25,000	,700					
Gaze	Sphericity Assumed	2,343	1	2,343	1,886	,182	,070	1,886	,262
	Greenhouse-Geisser	2,343	1,000	2,343	1,886	,182	,070	1,886	,262
	Huynh-Feldt	2,343	1,000	2,343	1,886	,182	,070	1,886	,262
	Lower-bound	2,343	1,000	2,343	1,886	,182	,070	1,886	,262
Gaze * Subject Extraversion	Sphericity Assumed	1,777	2	,889	,715	,499	,054	1,430	,157
	Greenhouse-Geisser	1,777	2,000	,889	,715	,499	,054	1,430	,157
	Huynh-Feldt	1,777	2,000	,889	,715	,499	,054	1,430	,157
	Lower-bound	1,777	2,000	,889	,715	,499	,054	1,430	,157
Gaze * Subject Agreeableness	Sphericity Assumed	1,117	2	,558	,449	,643	,035	,898	,115
	Greenhouse-Geisser	1,117	2,000	,558	,449	,643	,035	,898	,115
	Huynh-Feldt	1,117	2,000	,558	,449	,643	,035	,898	,115
	Lower-bound	1,117	2,000	,558	,449	,643	,035	,898	,115
Gaze * Subject Neuroticism	Sphericity Assumed	1,079	2	,539	,434	,653	,034	,868	,113
	Greenhouse-Geisser	1,079	2,000	,539	,434	,653	,034	,868	,113
	Huynh-Feldt	1,079	2,000	,539	,434	,653	,034	,868	,113
	Lower-bound	1,079	2,000	,539	,434	,653	,034	,868	,113
Error(Gaze)	Sphericity Assumed	31,068	25	1,243					
	Greenhouse-Geisser	31,068	25,000	1,243					
	Huynh-Feldt	31,068	25,000	1,243					
	Lower-bound	31,068	25,000	1,243					
Proxemics	Sphericity Assumed	117,607	1	117,607	67,197	,000	,729	67,197	1,000
	Greenhouse-Geisser	117,607	1,000	117,607	67,197	,000	,729	67,197	1,000
	Huynh-Feldt	117,607	1,000	117,607	67,197	,000	,729	67,197	1,000
	Lower-bound	117,607	1,000	117,607	67,197	,000	,729	67,197	1,000
Proxemics * Subject Extraversion	Sphericity Assumed	4,996	2	2,498	1,427	,259	,102	2,855	,277
	Greenhouse-Geisser	4,996	2,000	2,498	1,427	,259	,102	2,855	,277
	Huynh-Feldt	4,996	2,000	2,498	1,427	,259	,102	2,855	,277
	Lower-bound	4,996	2,000	2,498	1,427	,259	,102	2,855	,277
Proximity * Subject Agreeableness	Sphericity Assumed	,591	2	,296	,169	,846	,013	,338	,073
	Greenhouse-Geisser	,591	2,000	,296	,169	,846	,013	,338	,073
	Huynh-Feldt	,591	2,000	,296	,169	,846	,013	,338	,073
	Lower-bound	,591	2,000	,296	,169	,846	,013	,338	,073
Proximity * Subject Neuroticism	Sphericity Assumed	4,832	2	2,416	1,380	,270	,099	2,761	,269
	Greenhouse-Geisser	4,832	2,000	2,416	1,380	,270	,099	2,761	,269
	Huynh-Feldt	4,832	2,000	2,416	1,380	,270	,099	2,761	,269
	Lower-bound	4,832	2,000	2,416	1,380	,270	,099	2,761	,269
Error(Proxemics)	Sphericity Assumed	43,755	25	1,750					
	Greenhouse-Geisser	43,755	25,000	1,750					

(cont.)

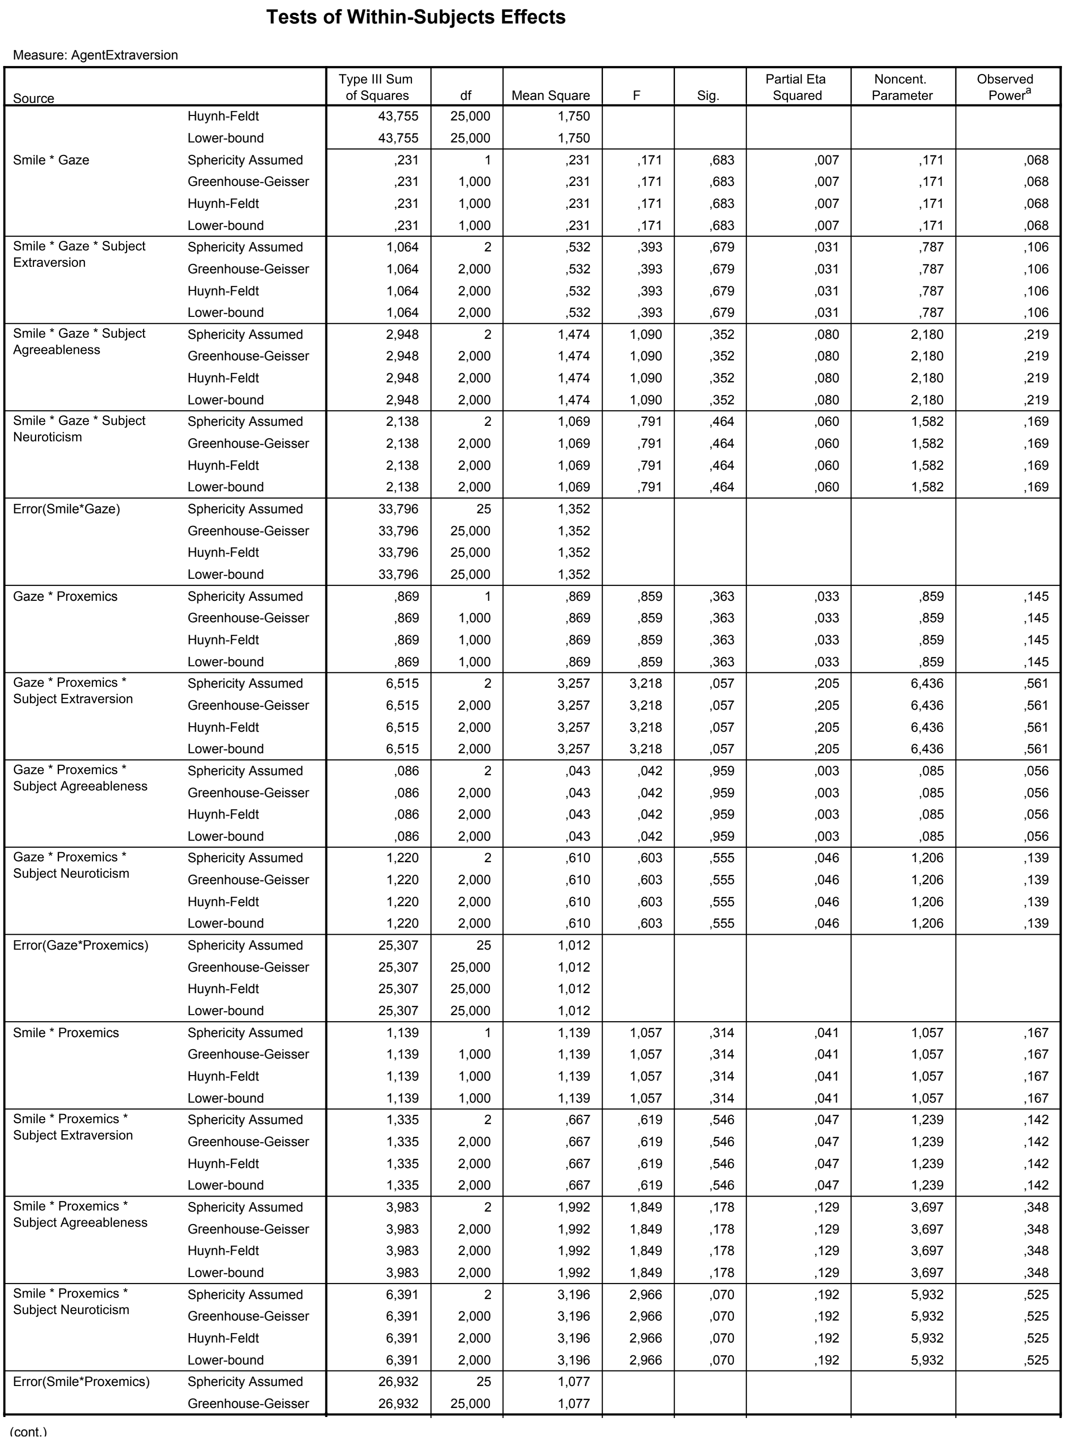

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power^a
	Huynh-Feldt	43,755	25,000	1,750					
	Lower-bound	43,755	25,000	1,750					
Smile * Gaze	Sphericity Assumed	,231	1	,231	,171	,683	,007	,171	,068
	Greenhouse-Geisser	,231	1,000	,231	,171	,683	,007	,171	,068
	Huynh-Feldt	,231	1,000	,231	,171	,683	,007	,171	,068
	Lower-bound	,231	1,000	,231	,171	,683	,007	,171	,068
Smile * Gaze * Subject Extraversion	Sphericity Assumed	1,064	2	,532	,393	,679	,031	,787	,106
	Greenhouse-Geisser	1,064	2,000	,532	,393	,679	,031	,787	,106
	Huynh-Feldt	1,064	2,000	,532	,393	,679	,031	,787	,106
	Lower-bound	1,064	2,000	,532	,393	,679	,031	,787	,106
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	2,948	2	1,474	1,090	,352	,080	2,180	,219
	Greenhouse-Geisser	2,948	2,000	1,474	1,090	,352	,080	2,180	,219
	Huynh-Feldt	2,948	2,000	1,474	1,090	,352	,080	2,180	,219
	Lower-bound	2,948	2,000	1,474	1,090	,352	,080	2,180	,219
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	2,138	2	1,069	,791	,464	,060	1,582	,169
	Greenhouse-Geisser	2,138	2,000	1,069	,791	,464	,060	1,582	,169
	Huynh-Feldt	2,138	2,000	1,069	,791	,464	,060	1,582	,169
	Lower-bound	2,138	2,000	1,069	,791	,464	,060	1,582	,169
Error(Smile*Gaze)	Sphericity Assumed	33,796	25	1,352					
	Greenhouse-Geisser	33,796	25,000	1,352					
	Huynh-Feldt	33,796	25,000	1,352					
	Lower-bound	33,796	25,000	1,352					
Gaze * Proxemics	Sphericity Assumed	,869	1	,869	,859	,363	,033	,859	,145
	Greenhouse-Geisser	,869	1,000	,869	,859	,363	,033	,859	,145
	Huynh-Feldt	,869	1,000	,869	,859	,363	,033	,859	,145
	Lower-bound	,869	1,000	,869	,859	,363	,033	,859	,145
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	6,515	2	3,257	3,218	,057	,205	6,436	,561
	Greenhouse-Geisser	6,515	2,000	3,257	3,218	,057	,205	6,436	,561
	Huynh-Feldt	6,515	2,000	3,257	3,218	,057	,205	6,436	,561
	Lower-bound	6,515	2,000	3,257	3,218	,057	,205	6,436	,561
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	,086	2	,043	,042	,959	,003	,085	,056
	Greenhouse-Geisser	,086	2,000	,043	,042	,959	,003	,085	,056
	Huynh-Feldt	,086	2,000	,043	,042	,959	,003	,085	,056
	Lower-bound	,086	2,000	,043	,042	,959	,003	,085	,056
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	1,220	2	,610	,603	,555	,046	1,206	,139
	Greenhouse-Geisser	1,220	2,000	,610	,603	,555	,046	1,206	,139
	Huynh-Feldt	1,220	2,000	,610	,603	,555	,046	1,206	,139
	Lower-bound	1,220	2,000	,610	,603	,555	,046	1,206	,139
Error(Gaze*Proxemics)	Sphericity Assumed	25,307	25	1,012					
	Greenhouse-Geisser	25,307	25,000	1,012					
	Huynh-Feldt	25,307	25,000	1,012					
	Lower-bound	25,307	25,000	1,012					
Smile * Proxemics	Sphericity Assumed	1,139	1	1,139	1,057	,314	,041	1,057	,167
	Greenhouse-Geisser	1,139	1,000	1,139	1,057	,314	,041	1,057	,167
	Huynh-Feldt	1,139	1,000	1,139	1,057	,314	,041	1,057	,167
	Lower-bound	1,139	1,000	1,139	1,057	,314	,041	1,057	,167
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	1,335	2	,667	,619	,546	,047	1,239	,142
	Greenhouse-Geisser	1,335	2,000	,667	,619	,546	,047	1,239	,142
	Huynh-Feldt	1,335	2,000	,667	,619	,546	,047	1,239	,142
	Lower-bound	1,335	2,000	,667	,619	,546	,047	1,239	,142
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	3,983	2	1,992	1,849	,178	,129	3,697	,348
	Greenhouse-Geisser	3,983	2,000	1,992	1,849	,178	,129	3,697	,348
	Huynh-Feldt	3,983	2,000	1,992	1,849	,178	,129	3,697	,348
	Lower-bound	3,983	2,000	1,992	1,849	,178	,129	3,697	,348
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	6,391	2	3,196	2,966	,070	,192	5,932	,525
	Greenhouse-Geisser	6,391	2,000	3,196	2,966	,070	,192	5,932	,525
	Huynh-Feldt	6,391	2,000	3,196	2,966	,070	,192	5,932	,525
	Lower-bound	6,391	2,000	3,196	2,966	,070	,192	5,932	,525
Error(Smile*Proxemics)	Sphericity Assumed	26,932	25	1,077					
	Greenhouse-Geisser	26,932	25,000	1,077					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentExtraversion

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	26,932	25,000	1,077					
	Lower-bound	26,932	25,000	1,077					
Smile * Gaze * Proxemics	Sphericity Assumed	,299	1	,299	,357	,555	,014	,357	,089
	Greenhouse-Geisser	,299	1,000	,299	,357	,555	,014	,357	,089
	Huynh-Feldt	,299	1,000	,299	,357	,555	,014	,357	,089
	Lower-bound	,299	1,000	,299	,357	,555	,014	,357	,089
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	,720	2	,360	,431	,655	,033	,861	,112
	Greenhouse-Geisser	,720	2,000	,360	,431	,655	,033	,861	,112
	Huynh-Feldt	,720	2,000	,360	,431	,655	,033	,861	,112
	Lower-bound	,720	2,000	,360	,431	,655	,033	,861	,112
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	,499	2	,249	,298	,745	,023	,596	,092
	Greenhouse-Geisser	,499	2,000	,249	,298	,745	,023	,596	,092
	Huynh-Feldt	,499	2,000	,249	,298	,745	,023	,596	,092
	Lower-bound	,499	2,000	,249	,298	,745	,023	,596	,092
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	2,161	2	1,080	1,291	,293	,094	2,583	,254
	Greenhouse-Geisser	2,161	2,000	1,080	1,291	,293	,094	2,583	,254
	Huynh-Feldt	2,161	2,000	1,080	1,291	,293	,094	2,583	,254
	Lower-bound	2,161	2,000	1,080	1,291	,293	,094	2,583	,254
Error (Smile*Gaze*Proxemics)	Sphericity Assumed	20,916	25	,837					
	Greenhouse-Geisser	20,916	25,000	,837					
	Huynh-Feldt	20,916	25,000	,837					
	Lower-bound	20,916	25,000	,837					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentExtraversion

Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound
5,533	,121	5,284	5,782

Smile

Estimates

Measure: AgentExtraversion

Smile	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,500	,130	5,232	5,768
2	5,566	,134	5,290	5,842

Pairwise Comparisons

Measure: AgentExtraversion

(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
					Lower Bound	Upper Bound
1	2	-,066	,107	,542	-,288	,155
2	1	,066	,107	,542	-,155	,288

Based on estimated marginal means

a. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,015	,383[a]	1,000	25,000	,542	,015	,383	,091
Wilks' lambda	,985	,383[a]	1,000	25,000	,542	,015	,383	,091
Hotelling's trace	,015	,383[a]	1,000	25,000	,542	,015	,383	,091
Roy's largest root	,015	,383[a]	1,000	25,000	,542	,015	,383	,091

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Proxemics

Estimates

Measure: AgentExtraversion

Proxemics	Mean	Std. Error	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
1	4,837	,148	4,532	5,143
2	6,229	,147	5,927	6,531

Pairwise Comparisons

Measure: AgentExtraversion

(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b] Lower Bound	95% Confidence Interval for Difference[b] Upper Bound
1	2	-1,391[*]	,170	,000	-1,741	-1,042
2	1	1,391[*]	,170	,000	1,042	1,741

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,729	67,197[a]	1,000	25,000	,000	,729	67,197	1,000
Wilks' lambda	,271	67,197[a]	1,000	25,000	,000	,729	67,197	1,000
Hotelling's trace	2,688	67,197[a]	1,000	25,000	,000	,729	67,197	1,000
Roy's largest root	2,688	67,197[a]	1,000	25,000	,000	,729	67,197	1,000

Each F tests the multivariate effect of Proxemics. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Estimates

Measure: AgentExtraversion

Smile	Subject Extraversion	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	5,379	,231	4,903	5,855
	MEDIUM	5,579	,233	5,100	6,059
	HIGH	5,541	,230	5,068	6,015
2	LOW	5,842	,238	5,352	6,333
	MEDIUM	5,243	,240	4,749	5,737
	HIGH	5,613	,237	5,125	6,101

Pairwise Comparisons

Measure: AgentExtraversion

Subject Extraversion	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
						Lower Bound	Upper Bound
LOW	1	2	-,463*	,191	,023	-,856	-,070
	2	1	,463*	,191	,023	,070	,856
MEDIUM	1	2	,336	,192	,093	-,060	,731
	2	1	-,336	,192	,093	-,731	,060
HIGH	1	2	-,072	,190	,708	-,462	,319
	2	1	,072	,190	,708	-,319	,462

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Extraversion		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,191	5,894[a]	1,000	25,000	,023	,191	5,894	,646
	Wilks' lambda	,809	5,894[a]	1,000	25,000	,023	,191	5,894	,646
	Hotelling's trace	,236	5,894[a]	1,000	25,000	,023	,191	5,894	,646
	Roy's largest root	,236	5,894[a]	1,000	25,000	,023	,191	5,894	,646
MEDIUM	Pillai's trace	,109	3,054[a]	1,000	25,000	,093	,109	3,054	,390
	Wilks' lambda	,891	3,054[a]	1,000	25,000	,093	,109	3,054	,390
	Hotelling's trace	,122	3,054[a]	1,000	25,000	,093	,109	3,054	,390
	Roy's largest root	,122	3,054[a]	1,000	25,000	,093	,109	3,054	,390
HIGH	Pillai's trace	,006	,144[a]	1,000	25,000	,708	,006	,144	,065
	Wilks' lambda	,994	,144[a]	1,000	25,000	,708	,006	,144	,065
	Hotelling's trace	,006	,144[a]	1,000	25,000	,708	,006	,144	,065
	Roy's largest root	,006	,144[a]	1,000	25,000	,708	,006	,144	,065

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Subject Neuroticism Factor * Gaze * Proxemics

Estimates

Measure: AgentExtraversion

					95% Confidence Interval	
Subject Neuroticism	Gaze	Proxemics	Mean	Std. Error	Lower Bound	Upper Bound
LOW	1	1	4,463	,398	3,645	5,282
		2	6,497	,321	5,836	7,158
	2	1	4,808	,349	4,089	5,527
		2	6,148	,376	5,374	6,923
MEDIUM	1	1	4,835	,317	4,182	5,488
		2	5,877	,256	5,350	6,404
	2	1	5,225	,278	4,651	5,798
		2	6,189	,300	5,571	6,806
HIGH	1	1	4,740	,320	4,080	5,399
		2	6,197	,259	5,665	6,730
	2	1	4,953	,281	4,374	5,532
		2	6,464	,303	5,841	7,088

Pairwise Comparisons

Measure: AgentExtraversion

							95% Confidence Interval for Difference[a]	
Subject Neuroticism	Proxemics	(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[a]	Lower Bound	Upper Bound
LOW	1	1	2	-,345	,433	,434	-1,237	,548
		2	1	,345	,433	,434	-,548	1,237
	2	1	2	,348	,354	,335	-,382	1,078
		2	1	-,348	,354	,335	-1,078	,382
MEDIUM	1	1	2	-,390	,346	,270	-1,102	,322
		2	1	,390	,346	,270	-,322	1,102
	2	1	2	-,311	,283	,281	-,894	,271
		2	1	,311	,283	,281	-,271	,894
HIGH	1	1	2	-,213	,349	,546	-,932	,505
		2	1	,213	,349	,546	-,505	,932
	2	1	2	-,267	,285	,359	-,855	,321
		2	1	,267	,285	,359	-,321	,855

Based on estimated marginal means

a. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Neuroticism	Proxemics		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	1	Pillai's trace	,025	,633[a]	1,000	25,000	,434	,025	,633	,119
		Wilks' lambda	,975	,633[a]	1,000	25,000	,434	,025	,633	,119
		Hotelling's trace	,025	,633[a]	1,000	25,000	,434	,025	,633	,119
		Roy's largest root	,025	,633[a]	1,000	25,000	,434	,025	,633	,119
	2	Pillai's trace	,037	,965[a]	1,000	25,000	,335	,037	,965	,157
		Wilks' lambda	,963	,965[a]	1,000	25,000	,335	,037	,965	,157
		Hotelling's trace	,039	,965[a]	1,000	25,000	,335	,037	,965	,157
		Roy's largest root	,039	,965[a]	1,000	25,000	,335	,037	,965	,157
MEDIUM	1	Pillai's trace	,048	1,272[a]	1,000	25,000	,270	,048	1,272	,192
		Wilks' lambda	,952	1,272[a]	1,000	25,000	,270	,048	1,272	,192
		Hotelling's trace	,051	1,272[a]	1,000	25,000	,270	,048	1,272	,192
		Roy's largest root	,051	1,272[a]	1,000	25,000	,270	,048	1,272	,192
	2	Pillai's trace	,046	1,213[a]	1,000	25,000	,281	,046	1,213	,185
		Wilks' lambda	,954	1,213[a]	1,000	25,000	,281	,046	1,213	,185
		Hotelling's trace	,049	1,213[a]	1,000	25,000	,281	,046	1,213	,185
		Roy's largest root	,049	1,213[a]	1,000	25,000	,281	,046	1,213	,185
HIGH	1	Pillai's trace	,015	,374[a]	1,000	25,000	,546	,015	,374	,090
		Wilks' lambda	,985	,374[a]	1,000	25,000	,546	,015	,374	,090
		Hotelling's trace	,015	,374[a]	1,000	25,000	,546	,015	,374	,090
		Roy's largest root	,015	,374[a]	1,000	25,000	,546	,015	,374	,090
	2	Pillai's trace	,034	,875[a]	1,000	25,000	,359	,034	,875	,147
		Wilks' lambda	,966	,875[a]	1,000	25,000	,359	,034	,875	,147
		Hotelling's trace	,035	,875[a]	1,000	25,000	,359	,034	,875	,147
		Roy's largest root	,035	,875[a]	1,000	25,000	,359	,034	,875	,147

Each F tests the multivariate simple effects of Gaze within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

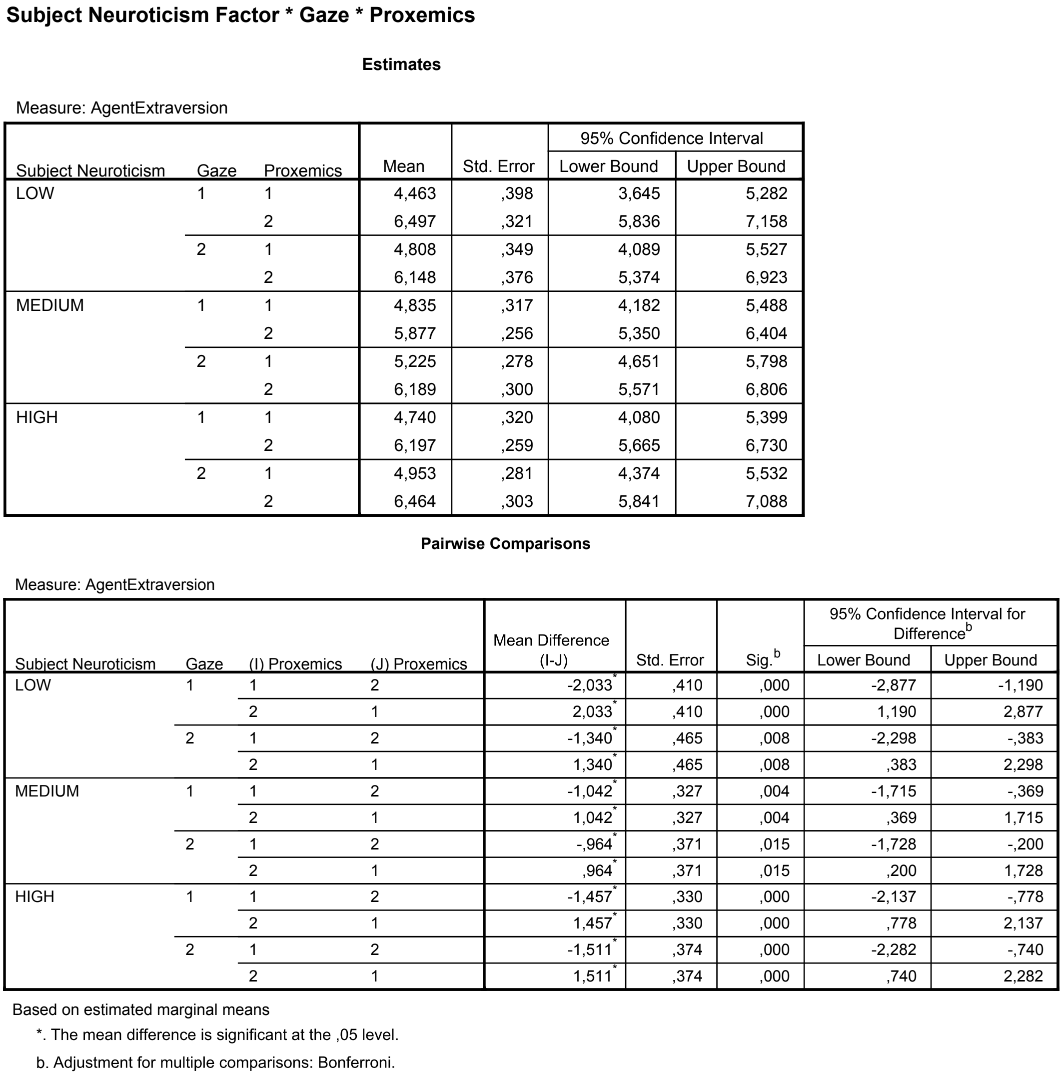

Subject Neuroticism Factor * Gaze * Proxemics

Estimates

Measure: AgentExtraversion

Subject Neuroticism	Gaze	Proxemics	Mean	Std. Error	95% Confidence Interval	
					Lower Bound	Upper Bound
LOW	1	1	4,463	,398	3,645	5,282
		2	6,497	,321	5,836	7,158
	2	1	4,808	,349	4,089	5,527
		2	6,148	,376	5,374	6,923
MEDIUM	1	1	4,835	,317	4,182	5,488
		2	5,877	,256	5,350	6,404
	2	1	5,225	,278	4,651	5,798
		2	6,189	,300	5,571	6,806
HIGH	1	1	4,740	,320	4,080	5,399
		2	6,197	,259	5,665	6,730
	2	1	4,953	,281	4,374	5,532
		2	6,464	,303	5,841	7,088

Pairwise Comparisons

Measure: AgentExtraversion

Subject Neuroticism	Gaze	(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
							Lower Bound	Upper Bound
LOW	1	1	2	-2,033[*]	,410	,000	-2,877	-1,190
		2	1	2,033[*]	,410	,000	1,190	2,877
	2	1	2	-1,340[*]	,465	,008	-2,298	-,383
		2	1	1,340[*]	,465	,008	,383	2,298
MEDIUM	1	1	2	-1,042[*]	,327	,004	-1,715	-,369
		2	1	1,042[*]	,327	,004	,369	1,715
	2	1	2	-,964[*]	,371	,015	-1,728	-,200
		2	1	,964[*]	,371	,015	,200	1,728
HIGH	1	1	2	-1,457[*]	,330	,000	-2,137	-,778
		2	1	1,457[*]	,330	,000	,778	2,137
	2	1	2	-1,511[*]	,374	,000	-2,282	-,740
		2	1	1,511[*]	,374	,000	,740	2,282

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Neuroticism	Gaze		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	1	Pillai's trace	,496	24,639[a]	1,000	25,000	,000	,496	24,639	,997
		Wilks' lambda	,504	24,639[a]	1,000	25,000	,000	,496	24,639	,997
		Hotelling's trace	,986	24,639[a]	1,000	25,000	,000	,496	24,639	,997
		Roy's largest root	,986	24,639[a]	1,000	25,000	,000	,496	24,639	,997
	2	Pillai's trace	,250	8,315[a]	1,000	25,000	,008	,250	8,315	,791
		Wilks' lambda	,750	8,315[a]	1,000	25,000	,008	,250	8,315	,791
		Hotelling's trace	,333	8,315[a]	1,000	25,000	,008	,250	8,315	,791
		Roy's largest root	,333	8,315[a]	1,000	25,000	,008	,250	8,315	,791
MEDIUM	1	Pillai's trace	,289	10,174[a]	1,000	25,000	,004	,289	10,174	,865
		Wilks' lambda	,711	10,174[a]	1,000	25,000	,004	,289	10,174	,865
		Hotelling's trace	,407	10,174[a]	1,000	25,000	,004	,289	10,174	,865
		Roy's largest root	,407	10,174[a]	1,000	25,000	,004	,289	10,174	,865
	2	Pillai's trace	,213	6,755[a]	1,000	25,000	,015	,213	6,755	,705
		Wilks' lambda	,787	6,755[a]	1,000	25,000	,015	,213	6,755	,705
		Hotelling's trace	,270	6,755[a]	1,000	25,000	,015	,213	6,755	,705
		Roy's largest root	,270	6,755[a]	1,000	25,000	,015	,213	6,755	,705
HIGH	1	Pillai's trace	,438	19,512[a]	1,000	25,000	,000	,438	19,512	,989
		Wilks' lambda	,562	19,512[a]	1,000	25,000	,000	,438	19,512	,989
		Hotelling's trace	,780	19,512[a]	1,000	25,000	,000	,438	19,512	,989
		Roy's largest root	,780	19,512[a]	1,000	25,000	,000	,438	19,512	,989
	2	Pillai's trace	,394	16,287[a]	1,000	25,000	,000	,394	16,287	,972
		Wilks' lambda	,606	16,287[a]	1,000	25,000	,000	,394	16,287	,972
		Hotelling's trace	,651	16,287[a]	1,000	25,000	,000	,394	16,287	,972
		Roy's largest root	,651	16,287[a]	1,000	25,000	,000	,394	16,287	,972

Each F tests the multivariate simple effects of Proxemics within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Agent Friendliness Measure

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Smile	Sphericity Assumed	89,618	1	89,618	49,074	,000	,662	49,074	1,000
	Greenhouse-Geisser	89,618	1,000	89,618	49,074	,000	,662	49,074	1,000
	Huynh-Feldt	89,618	1,000	89,618	49,074	,000	,662	49,074	1,000
	Lower-bound	89,618	1,000	89,618	49,074	,000	,662	49,074	1,000
Smile * Subject Extraversion	Sphericity Assumed	29,232	2	14,616	8,004	,002	,390	16,007	,931
	Greenhouse-Geisser	29,232	2,000	14,616	8,004	,002	,390	16,007	,931
	Huynh-Feldt	29,232	2,000	14,616	8,004	,002	,390	16,007	,931
	Lower-bound	29,232	2,000	14,616	8,004	,002	,390	16,007	,931
Smile * Subject Agreeableness	Sphericity Assumed	11,263	2	5,632	3,084	,064	,198	6,168	,542
	Greenhouse-Geisser	11,263	2,000	5,632	3,084	,064	,198	6,168	,542
	Huynh-Feldt	11,263	2,000	5,632	3,084	,064	,198	6,168	,542
	Lower-bound	11,263	2,000	5,632	3,084	,064	,198	6,168	,542
Smile * Subject Neuroticism	Sphericity Assumed	,546	2	,273	,149	,862	,012	,299	,071
	Greenhouse-Geisser	,546	2,000	,273	,149	,862	,012	,299	,071
	Huynh-Feldt	,546	2,000	,273	,149	,862	,012	,299	,071
	Lower-bound	,546	2,000	,273	,149	,862	,012	,299	,071
Error(Smile)	Sphericity Assumed	45,654	25	1,826					
	Greenhouse-Geisser	45,654	25,000	1,826					
	Huynh-Feldt	45,654	25,000	1,826					
	Lower-bound	45,654	25,000	1,826					
Gaze	Sphericity Assumed	15,714	1	15,714	12,328	,002	,330	12,328	,921
	Greenhouse-Geisser	15,714	1,000	15,714	12,328	,002	,330	12,328	,921
	Huynh-Feldt	15,714	1,000	15,714	12,328	,002	,330	12,328	,921
	Lower-bound	15,714	1,000	15,714	12,328	,002	,330	12,328	,921
Gaze * Subject Extraversion	Sphericity Assumed	5,602	2	2,801	2,198	,132	,150	4,395	,406
	Greenhouse-Geisser	5,602	2,000	2,801	2,198	,132	,150	4,395	,406
	Huynh-Feldt	5,602	2,000	2,801	2,198	,132	,150	4,395	,406
	Lower-bound	5,602	2,000	2,801	2,198	,132	,150	4,395	,406
Gaze * Subject Agreeableness	Sphericity Assumed	5,936	2	2,968	2,329	,118	,157	4,657	,427
	Greenhouse-Geisser	5,936	2,000	2,968	2,329	,118	,157	4,657	,427
	Huynh-Feldt	5,936	2,000	2,968	2,329	,118	,157	4,657	,427
	Lower-bound	5,936	2,000	2,968	2,329	,118	,157	4,657	,427
Gaze * Subject Neuroticism	Sphericity Assumed	6,366	2	3,183	2,497	,103	,167	4,995	,454
	Greenhouse-Geisser	6,366	2,000	3,183	2,497	,103	,167	4,995	,454
	Huynh-Feldt	6,366	2,000	3,183	2,497	,103	,167	4,995	,454
	Lower-bound	6,366	2,000	3,183	2,497	,103	,167	4,995	,454
Error(Gaze)	Sphericity Assumed	31,865	25	1,275					
	Greenhouse-Geisser	31,865	25,000	1,275					
	Huynh-Feldt	31,865	25,000	1,275					
	Lower-bound	31,865	25,000	1,275					
Proxemics	Sphericity Assumed	3,775	1	3,775	1,104	,303	,042	1,104	,173
	Greenhouse-Geisser	3,775	1,000	3,775	1,104	,303	,042	1,104	,173
	Huynh-Feldt	3,775	1,000	3,775	1,104	,303	,042	1,104	,173
	Lower-bound	3,775	1,000	3,775	1,104	,303	,042	1,104	,173
Proxemics * Subject Extraversion	Sphericity Assumed	11,751	2	5,875	1,719	,200	,121	3,437	,326
	Greenhouse-Geisser	11,751	2,000	5,875	1,719	,200	,121	3,437	,326
	Huynh-Feldt	11,751	2,000	5,875	1,719	,200	,121	3,437	,326
	Lower-bound	11,751	2,000	5,875	1,719	,200	,121	3,437	,326
Proxemics * Subject Agreeableness	Sphericity Assumed	15,890	2	7,945	2,324	,119	,157	4,648	,427
	Greenhouse-Geisser	15,890	2,000	7,945	2,324	,119	,157	4,648	,427
	Huynh-Feldt	15,890	2,000	7,945	2,324	,119	,157	4,648	,427
	Lower-bound	15,890	2,000	7,945	2,324	,119	,157	4,648	,427
Proxemics * Subject Neuroticism	Sphericity Assumed	17,334	2	8,667	2,535	,099	,169	5,071	,460
	Greenhouse-Geisser	17,334	2,000	8,667	2,535	,099	,169	5,071	,460
	Huynh-Feldt	17,334	2,000	8,667	2,535	,099	,169	5,071	,460
	Lower-bound	17,334	2,000	8,667	2,535	,099	,169	5,071	,460
Error(Proxemics)	Sphericity Assumed	85,464	25	3,419					
	Greenhouse-Geisser	85,464	25,000	3,419					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	85,464	25,000	3,419					
	Lower-bound	85,464	25,000	3,419					
Smile * Gaze	Sphericity Assumed	,500	1	,500	,255	,618	,010	,255	,077
	Greenhouse-Geisser	,500	1,000	,500	,255	,618	,010	,255	,077
	Huynh-Feldt	,500	1,000	,500	,255	,618	,010	,255	,077
	Lower-bound	,500	1,000	,500	,255	,618	,010	,255	,077
Smile * Gaze * Subject Extraversion	Sphericity Assumed	2,295	2	1,148	,586	,564	,045	1,173	,136
	Greenhouse-Geisser	2,295	2,000	1,148	,586	,564	,045	1,173	,136
	Huynh-Feldt	2,295	2,000	1,148	,586	,564	,045	1,173	,136
	Lower-bound	2,295	2,000	1,148	,586	,564	,045	1,173	,136
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	1,211	2	,605	,309	,737	,024	,619	,094
	Greenhouse-Geisser	1,211	2,000	,605	,309	,737	,024	,619	,094
	Huynh-Feldt	1,211	2,000	,605	,309	,737	,024	,619	,094
	Lower-bound	1,211	2,000	,605	,309	,737	,024	,619	,094
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	8,857	2	4,429	2,263	,125	,153	4,526	,417
	Greenhouse-Geisser	8,857	2,000	4,429	2,263	,125	,153	4,526	,417
	Huynh-Feldt	8,857	2,000	4,429	2,263	,125	,153	4,526	,417
	Lower-bound	8,857	2,000	4,429	2,263	,125	,153	4,526	,417
Error(Smile*Gaze)	Sphericity Assumed	48,925	25	1,957					
	Greenhouse-Geisser	48,925	25,000	1,957					
	Huynh-Feldt	48,925	25,000	1,957					
	Lower-bound	48,925	25,000	1,957					
Smile * Proxemics	Sphericity Assumed	,001	1	,001	,000	,985	,000	,000	,050
	Greenhouse-Geisser	,001	1,000	,001	,000	,985	,000	,000	,050
	Huynh-Feldt	,001	1,000	,001	,000	,985	,000	,000	,050
	Lower-bound	,001	1,000	,001	,000	,985	,000	,000	,050
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	6,039	2	3,020	1,167	,328	,085	2,335	,232
	Greenhouse-Geisser	6,039	2,000	3,020	1,167	,328	,085	2,335	,232
	Huynh-Feldt	6,039	2,000	3,020	1,167	,328	,085	2,335	,232
	Lower-bound	6,039	2,000	3,020	1,167	,328	,085	2,335	,232
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	8,282	2	4,141	1,601	,222	,114	3,202	,306
	Greenhouse-Geisser	8,282	2,000	4,141	1,601	,222	,114	3,202	,306
	Huynh-Feldt	8,282	2,000	4,141	1,601	,222	,114	3,202	,306
	Lower-bound	8,282	2,000	4,141	1,601	,222	,114	3,202	,306
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	1,573	2	,787	,304	,740	,024	,608	,093
	Greenhouse-Geisser	1,573	2,000	,787	,304	,740	,024	,608	,093
	Huynh-Feldt	1,573	2,000	,787	,304	,740	,024	,608	,093
	Lower-bound	1,573	2,000	,787	,304	,740	,024	,608	,093
Error(Smile*Proxemics)	Sphericity Assumed	64,661	25	2,586					
	Greenhouse-Geisser	64,661	25,000	2,586					
	Huynh-Feldt	64,661	25,000	2,586					
	Lower-bound	64,661	25,000	2,586					
Gaze * Proxemics	Sphericity Assumed	,827	1	,827	,605	,444	,024	,605	,116
	Greenhouse-Geisser	,827	1,000	,827	,605	,444	,024	,605	,116
	Huynh-Feldt	,827	1,000	,827	,605	,444	,024	,605	,116
	Lower-bound	,827	1,000	,827	,605	,444	,024	,605	,116
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	6,420	2	3,210	2,347	,116	,158	4,693	,430
	Greenhouse-Geisser	6,420	2,000	3,210	2,347	,116	,158	4,693	,430
	Huynh-Feldt	6,420	2,000	3,210	2,347	,116	,158	4,693	,430
	Lower-bound	6,420	2,000	3,210	2,347	,116	,158	4,693	,430
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	7,795	2	3,898	2,849	,077	,186	5,698	,508
	Greenhouse-Geisser	7,795	2,000	3,898	2,849	,077	,186	5,698	,508
	Huynh-Feldt	7,795	2,000	3,898	2,849	,077	,186	5,698	,508
	Lower-bound	7,795	2,000	3,898	2,849	,077	,186	5,698	,508
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	3,831	2	1,916	1,400	,265	,101	2,801	,272
	Greenhouse-Geisser	3,831	2,000	1,916	1,400	,265	,101	2,801	,272
	Huynh-Feldt	3,831	2,000	1,916	1,400	,265	,101	2,801	,272
	Lower-bound	3,831	2,000	1,916	1,400	,265	,101	2,801	,272
Error(Gaze*Proxemics)	Sphericity Assumed	34,199	25	1,368					
	Greenhouse-Geisser	34,199	25,000	1,368					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentFriendliness

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	34,199	25,000	1,368					
	Lower-bound	34,199	25,000	1,368					
Smile * Gaze * Proxemics	Sphericity Assumed	2,671	1	2,671	2,061	,163	,076	2,061	,282
	Greenhouse-Geisser	2,671	1,000	2,671	2,061	,163	,076	2,061	,282
	Huynh-Feldt	2,671	1,000	2,671	2,061	,163	,076	2,061	,282
	Lower-bound	2,671	1,000	2,671	2,061	,163	,076	2,061	,282
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	1,275	2	,637	,492	,617	,038	,984	,122
	Greenhouse-Geisser	1,275	2,000	,637	,492	,617	,038	,984	,122
	Huynh-Feldt	1,275	2,000	,637	,492	,617	,038	,984	,122
	Lower-bound	1,275	2,000	,637	,492	,617	,038	,984	,122
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	7,713	2	3,856	2,976	,069	,192	5,953	,527
	Greenhouse-Geisser	7,713	2,000	3,856	2,976	,069	,192	5,953	,527
	Huynh-Feldt	7,713	2,000	3,856	2,976	,069	,192	5,953	,527
	Lower-bound	7,713	2,000	3,856	2,976	,069	,192	5,953	,527
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,692	2	,346	,267	,768	,021	,534	,087
	Greenhouse-Geisser	,692	2,000	,346	,267	,768	,021	,534	,087
	Huynh-Feldt	,692	2,000	,346	,267	,768	,021	,534	,087
	Lower-bound	,692	2,000	,346	,267	,768	,021	,534	,087
Error (Smile*Gaze*Proxemics)	Sphericity Assumed	32,390	25	1,296					
	Greenhouse-Geisser	32,390	25,000	1,296					
	Huynh-Feldt	32,390	25,000	1,296					
	Lower-bound	32,390	25,000	1,296					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentFriendliness

Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound
5,866	,118	5,623	6,109

Smile

Estimates

Measure: AgentFriendliness

Smile	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,258	,134	4,983	5,534
2	6,473	,158	6,147	6,799

Pairwise Comparisons

Measure: AgentFriendliness

(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
					Lower Bound	Upper Bound
1	2	-1,215*	,173	,000	-1,572	-,857
2	1	1,215*	,173	,000	,857	1,572

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,662	49,074[a]	1,000	25,000	,000	,662	49,074	1,000
Wilks' lambda	,338	49,074[a]	1,000	25,000	,000	,662	49,074	1,000
Hotelling's trace	1,963	49,074[a]	1,000	25,000	,000	,662	49,074	1,000
Roy's largest root	1,963	49,074[a]	1,000	25,000	,000	,662	49,074	1,000

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze

Estimates

Measure: AgentFriendliness

Gaze	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	5,611	,152	5,299	5,924
2	6,120	,123	5,866	6,374

Pairwise Comparisons

Measure: AgentFriendliness

(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
					Lower Bound	Upper Bound
1	2	-,509[*]	,145	,002	-,807	-,210
2	1	,509[*]	,145	,002	,210	,807

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,330	12,328[a]	1,000	25,000	,002	,330	12,328	,921
Wilks' lambda	,670	12,328[a]	1,000	25,000	,002	,330	12,328	,921
Hotelling's trace	,493	12,328[a]	1,000	25,000	,002	,330	12,328	,921
Roy's largest root	,493	12,328[a]	1,000	25,000	,002	,330	12,328	,921

Each F tests the multivariate effect of Gaze. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Smile * Subject Extraversion Factor

Estimates

Measure: AgentFriendliness

Smile	Subject Extraversion	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	5,656	,238	5,167	6,146
	MEDIUM	4,682	,239	4,189	5,175
	HIGH	5,437	,236	4,951	5,924
2	LOW	6,151	,281	5,573	6,730
	MEDIUM	6,875	,283	6,292	7,457
	HIGH	6,393	,279	5,818	6,968

Pairwise Comparisons

Measure: AgentFriendliness

Subject Extraversion	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
						Lower Bound	Upper Bound
LOW	1	2	-,495	,308	,121	-1,130	,139
	2	1	,495	,308	,121	-,139	1,130
MEDIUM	1	2	-2,193*	,310	,000	-2,832	-1,554
	2	1	2,193*	,310	,000	1,554	2,832
HIGH	1	2	-,956*	,306	,005	-1,587	-,325
	2	1	,956*	,306	,005	,325	1,587

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Extraversion		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,094	2,582[a]	1,000	25,000	,121	,094	2,582	,339
	Wilks' lambda	,906	2,582[a]	1,000	25,000	,121	,094	2,582	,339
	Hotelling's trace	,103	2,582[a]	1,000	25,000	,121	,094	2,582	,339
	Roy's largest root	,103	2,582[a]	1,000	25,000	,121	,094	2,582	,339
MEDIUM	Pillai's trace	,666	49,941[a]	1,000	25,000	,000	,666	49,941	1,000
	Wilks' lambda	,334	49,941[a]	1,000	25,000	,000	,666	49,941	1,000
	Hotelling's trace	1,998	49,941[a]	1,000	25,000	,000	,666	49,941	1,000
	Roy's largest root	1,998	49,941[a]	1,000	25,000	,000	,666	49,941	1,000
HIGH	Pillai's trace	,280	9,741[a]	1,000	25,000	,005	,280	9,741	,851
	Wilks' lambda	,720	9,741[a]	1,000	25,000	,005	,280	9,741	,851
	Hotelling's trace	,390	9,741[a]	1,000	25,000	,005	,280	9,741	,851
	Roy's largest root	,390	9,741[a]	1,000	25,000	,005	,280	9,741	,851

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Smile * Subject Agreeableness Factor

Estimates

Measure: AgentFriendliness

Smile	Subject Agreeableness	Mean	Std. Error	95% Confidence Interval Lower Bound	95% Confidence Interval Upper Bound
1	LOW	4,603	,260	4,068	5,137
	MEDIUM	5,524	,217	5,078	5,971
	HIGH	5,649	,242	5,150	6,147
2	LOW	6,495	,307	5,863	7,127
	MEDIUM	6,420	,256	5,892	6,948
	HIGH	6,505	,286	5,915	7,094

Pairwise Comparisons

Measure: AgentFriendliness

Subject Agreeableness	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b] Lower Bound	95% Confidence Interval for Difference[b] Upper Bound
LOW	1	2	-1,892*	,337	,000	-2,585	-1,199
	2	1	1,892*	,337	,000	1,199	2,585
MEDIUM	1	2	-,896*	,281	,004	-1,475	-,317
	2	1	,896*	,281	,004	,317	1,475
HIGH	1	2	-,856*	,314	,012	-1,503	-,209
	2	1	,856*	,314	,012	,209	1,503

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Agreeableness		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,558	31,599[a]	1,000	25,000	,000	,558	31,599	1,000
	Wilks' lambda	,442	31,599[a]	1,000	25,000	,000	,558	31,599	1,000
	Hotelling's trace	1,264	31,599[a]	1,000	25,000	,000	,558	31,599	1,000
	Roy's largest root	1,264	31,599[a]	1,000	25,000	,000	,558	31,599	1,000
MEDIUM	Pillai's trace	,289	10,153[a]	1,000	25,000	,004	,289	10,153	,865
	Wilks' lambda	,711	10,153[a]	1,000	25,000	,004	,289	10,153	,865
	Hotelling's trace	,406	10,153[a]	1,000	25,000	,004	,289	10,153	,865
	Roy's largest root	,406	10,153[a]	1,000	25,000	,004	,289	10,153	,865
HIGH	Pillai's trace	,229	7,430[a]	1,000	25,000	,012	,229	7,430	,745
	Wilks' lambda	,771	7,430[a]	1,000	25,000	,012	,229	7,430	,745
	Hotelling's trace	,297	7,430[a]	1,000	25,000	,012	,229	7,430	,745
	Roy's largest root	,297	7,430[a]	1,000	25,000	,012	,229	7,430	,745

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Agent Likeability Measure

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Smile	Sphericity Assumed	36,305	1	36,305	41,349	,000	,623	41,349	1,000
	Greenhouse-Geisser	36,305	1,000	36,305	41,349	,000	,623	41,349	1,000
	Huynh-Feldt	36,305	1,000	36,305	41,349	,000	,623	41,349	1,000
	Lower-bound	36,305	1,000	36,305	41,349	,000	,623	41,349	1,000
Smile * Subject Extraversion	Sphericity Assumed	4,708	2	2,354	2,681	,088	,177	5,363	,483
	Greenhouse-Geisser	4,708	2,000	2,354	2,681	,088	,177	5,363	,483
	Huynh-Feldt	4,708	2,000	2,354	2,681	,088	,177	5,363	,483
	Lower-bound	4,708	2,000	2,354	2,681	,088	,177	5,363	,483
Smile * Subject Agreeableness	Sphericity Assumed	4,022	2	2,011	2,291	,122	,155	4,581	,421
	Greenhouse-Geisser	4,022	2,000	2,011	2,291	,122	,155	4,581	,421
	Huynh-Feldt	4,022	2,000	2,011	2,291	,122	,155	4,581	,421
	Lower-bound	4,022	2,000	2,011	2,291	,122	,155	4,581	,421
Smile * Subject Neuroticism	Sphericity Assumed	,114	2	,057	,065	,937	,005	,130	,059
	Greenhouse-Geisser	,114	2,000	,057	,065	,937	,005	,130	,059
	Huynh-Feldt	,114	2,000	,057	,065	,937	,005	,130	,059
	Lower-bound	,114	2,000	,057	,065	,937	,005	,130	,059
Error(Smile)	Sphericity Assumed	21,950	25	,878					
	Greenhouse-Geisser	21,950	25,000	,878					
	Huynh-Feldt	21,950	25,000	,878					
	Lower-bound	21,950	25,000	,878					
Gaze	Sphericity Assumed	5,101	1	5,101	9,910	,004	,284	9,910	,857
	Greenhouse-Geisser	5,101	1,000	5,101	9,910	,004	,284	9,910	,857
	Huynh-Feldt	5,101	1,000	5,101	9,910	,004	,284	9,910	,857
	Lower-bound	5,101	1,000	5,101	9,910	,004	,284	9,910	,857
Gaze * Subject Extraversion	Sphericity Assumed	1,893	2	,947	1,839	,180	,128	3,678	,347
	Greenhouse-Geisser	1,893	2,000	,947	1,839	,180	,128	3,678	,347
	Huynh-Feldt	1,893	2,000	,947	1,839	,180	,128	3,678	,347
	Lower-bound	1,893	2,000	,947	1,839	,180	,128	3,678	,347
Gaze * Subject Agreeableness	Sphericity Assumed	1,693	2	,846	1,644	,213	,116	3,289	,314
	Greenhouse-Geisser	1,693	2,000	,846	1,644	,213	,116	3,289	,314
	Huynh-Feldt	1,693	2,000	,846	1,644	,213	,116	3,289	,314
	Lower-bound	1,693	2,000	,846	1,644	,213	,116	3,289	,314
Gaze * Subject Neuroticism	Sphericity Assumed	,007	2	,003	,006	,994	,001	,013	,051
	Greenhouse-Geisser	,007	2,000	,003	,006	,994	,001	,013	,051
	Huynh-Feldt	,007	2,000	,003	,006	,994	,001	,013	,051
	Lower-bound	,007	2,000	,003	,006	,994	,001	,013	,051
Error(Gaze)	Sphericity Assumed	12,869	25	,515					
	Greenhouse-Geisser	12,869	25,000	,515					
	Huynh-Feldt	12,869	25,000	,515					
	Lower-bound	12,869	25,000	,515					
Proxemics	Sphericity Assumed	,034	1	,034	,021	,886	,001	,021	,052
	Greenhouse-Geisser	,034	1,000	,034	,021	,886	,001	,021	,052
	Huynh-Feldt	,034	1,000	,034	,021	,886	,001	,021	,052
	Lower-bound	,034	1,000	,034	,021	,886	,001	,021	,052
Proxemics * Subject Extraversion	Sphericity Assumed	8,685	2	4,343	2,734	,084	,179	5,469	,491
	Greenhouse-Geisser	8,685	2,000	4,343	2,734	,084	,179	5,469	,491
	Huynh-Feldt	8,685	2,000	4,343	2,734	,084	,179	5,469	,491
	Lower-bound	8,685	2,000	4,343	2,734	,084	,179	5,469	,491
Proxemics * Subject Agreeableness	Sphericity Assumed	7,219	2	3,610	2,273	,124	,154	4,546	,418
	Greenhouse-Geisser	7,219	2,000	3,610	2,273	,124	,154	4,546	,418
	Huynh-Feldt	7,219	2,000	3,610	2,273	,124	,154	4,546	,418
	Lower-bound	7,219	2,000	3,610	2,273	,124	,154	4,546	,418
Proxemics * Subject Neuroticism	Sphericity Assumed	6,812	2	3,406	2,145	,138	,146	4,290	,397
	Greenhouse-Geisser	6,812	2,000	3,406	2,145	,138	,146	4,290	,397
	Huynh-Feldt	6,812	2,000	3,406	2,145	,138	,146	4,290	,397
	Lower-bound	6,812	2,000	3,406	2,145	,138	,146	4,290	,397
Error(Proxemics)	Sphericity Assumed	39,701	25	1,588					
	Greenhouse-Geisser	39,701	25,000	1,588					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	39,701	25,000	1,588					
	Lower-bound	39,701	25,000	1,588					
Smile * Gaze	Sphericity Assumed	1,061	1	1,061	,948	,340	,037	,948	,155
	Greenhouse-Geisser	1,061	1,000	1,061	,948	,340	,037	,948	,155
	Huynh-Feldt	1,061	1,000	1,061	,948	,340	,037	,948	,155
	Lower-bound	1,061	1,000	1,061	,948	,340	,037	,948	,155
Smile * Gaze * Subject Extraversion	Sphericity Assumed	1,073	2	,537	,480	,625	,037	,959	,120
	Greenhouse-Geisser	1,073	2,000	,537	,480	,625	,037	,959	,120
	Huynh-Feldt	1,073	2,000	,537	,480	,625	,037	,959	,120
	Lower-bound	1,073	2,000	,537	,480	,625	,037	,959	,120
Smile * Gaze * Subject Agreeableness	Sphericity Assumed	,274	2	,137	,123	,885	,010	,245	,067
	Greenhouse-Geisser	,274	2,000	,137	,123	,885	,010	,245	,067
	Huynh-Feldt	,274	2,000	,137	,123	,885	,010	,245	,067
	Lower-bound	,274	2,000	,137	,123	,885	,010	,245	,067
Smile * Gaze * Subject Neuroticism	Sphericity Assumed	,834	2	,417	,373	,693	,029	,746	,103
	Greenhouse-Geisser	,834	2,000	,417	,373	,693	,029	,746	,103
	Huynh-Feldt	,834	2,000	,417	,373	,693	,029	,746	,103
	Lower-bound	,834	2,000	,417	,373	,693	,029	,746	,103
Error(Smile*Gaze)	Sphericity Assumed	27,974	25	1,119					
	Greenhouse-Geisser	27,974	25,000	1,119					
	Huynh-Feldt	27,974	25,000	1,119					
	Lower-bound	27,974	25,000	1,119					
Smile * Proxemics	Sphericity Assumed	,179	1	,179	,218	,644	,009	,218	,073
	Greenhouse-Geisser	,179	1,000	,179	,218	,644	,009	,218	,073
	Huynh-Feldt	,179	1,000	,179	,218	,644	,009	,218	,073
	Lower-bound	,179	1,000	,179	,218	,644	,009	,218	,073
Smile * Proxemics * Subject Extraversion	Sphericity Assumed	2,265	2	1,132	1,378	,271	,099	2,755	,268
	Greenhouse-Geisser	2,265	2,000	1,132	1,378	,271	,099	2,755	,268
	Huynh-Feldt	2,265	2,000	1,132	1,378	,271	,099	2,755	,268
	Lower-bound	2,265	2,000	1,132	1,378	,271	,099	2,755	,268
Smile * Proxemics * Subject Agreeableness	Sphericity Assumed	2,106	2	1,053	1,281	,295	,093	2,562	,252
	Greenhouse-Geisser	2,106	2,000	1,053	1,281	,295	,093	2,562	,252
	Huynh-Feldt	2,106	2,000	1,053	1,281	,295	,093	2,562	,252
	Lower-bound	2,106	2,000	1,053	1,281	,295	,093	2,562	,252
Smile * Proxemics * Subject Neuroticism	Sphericity Assumed	,670	2	,335	,408	,669	,032	,816	,109
	Greenhouse-Geisser	,670	2,000	,335	,408	,669	,032	,816	,109
	Huynh-Feldt	,670	2,000	,335	,408	,669	,032	,816	,109
	Lower-bound	,670	2,000	,335	,408	,669	,032	,816	,109
Error(Smile*Proxemics)	Sphericity Assumed	20,546	25	,822					
	Greenhouse-Geisser	20,546	25,000	,822					
	Huynh-Feldt	20,546	25,000	,822					
	Lower-bound	20,546	25,000	,822					
Gaze * Proxemics	Sphericity Assumed	,138	1	,138	,156	,697	,006	,156	,067
	Greenhouse-Geisser	,138	1,000	,138	,156	,697	,006	,156	,067
	Huynh-Feldt	,138	1,000	,138	,156	,697	,006	,156	,067
	Lower-bound	,138	1,000	,138	,156	,697	,006	,156	,067
Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	,103	2	,052	,058	,943	,005	,117	,058
	Greenhouse-Geisser	,103	2,000	,052	,058	,943	,005	,117	,058
	Huynh-Feldt	,103	2,000	,052	,058	,943	,005	,117	,058
	Lower-bound	,103	2,000	,052	,058	,943	,005	,117	,058
Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	3,357	2	1,678	1,897	,171	,132	3,795	,356
	Greenhouse-Geisser	3,357	2,000	1,678	1,897	,171	,132	3,795	,356
	Huynh-Feldt	3,357	2,000	1,678	1,897	,171	,132	3,795	,356
	Lower-bound	3,357	2,000	1,678	1,897	,171	,132	3,795	,356
Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,504	2	,252	,285	,755	,022	,570	,090
	Greenhouse-Geisser	,504	2,000	,252	,285	,755	,022	,570	,090
	Huynh-Feldt	,504	2,000	,252	,285	,755	,022	,570	,090
	Lower-bound	,504	2,000	,252	,285	,755	,022	,570	,090
Error(Gaze*Proxemics)	Sphericity Assumed	22,114	25	,885					
	Greenhouse-Geisser	22,114	25,000	,885					

(cont.)

Tests of Within-Subjects Effects

Measure: AgentLikeability

Source		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
	Huynh-Feldt	22,114	25,000	,885					
	Lower-bound	22,114	25,000	,885					
Smile * Gaze * Proxemics	Sphericity Assumed	2,406	1	2,406	2,843	,104	,102	2,843	,368
	Greenhouse-Geisser	2,406	1,000	2,406	2,843	,104	,102	2,843	,368
	Huynh-Feldt	2,406	1,000	2,406	2,843	,104	,102	2,843	,368
	Lower-bound	2,406	1,000	2,406	2,843	,104	,102	2,843	,368
Smile * Gaze * Proxemics * Subject Extraversion	Sphericity Assumed	,994	2	,497	,587	,563	,045	1,175	,137
	Greenhouse-Geisser	,994	2,000	,497	,587	,563	,045	1,175	,137
	Huynh-Feldt	,994	2,000	,497	,587	,563	,045	1,175	,137
	Lower-bound	,994	2,000	,497	,587	,563	,045	1,175	,137
Smile * Gaze * Proxemics * Subject Agreeableness	Sphericity Assumed	3,250	2	1,625	1,920	,168	,133	3,841	,360
	Greenhouse-Geisser	3,250	2,000	1,625	1,920	,168	,133	3,841	,360
	Huynh-Feldt	3,250	2,000	1,625	1,920	,168	,133	3,841	,360
	Lower-bound	3,250	2,000	1,625	1,920	,168	,133	3,841	,360
Smile * Gaze * Proxemics * Subject Neuroticism	Sphericity Assumed	,944	2	,472	,558	,579	,043	1,116	,132
	Greenhouse-Geisser	,944	2,000	,472	,558	,579	,043	1,116	,132
	Huynh-Feldt	,944	2,000	,472	,558	,579	,043	1,116	,132
	Lower-bound	,944	2,000	,472	,558	,579	,043	1,116	,132
Error (Smile*Gaze*Proxemics)	Sphericity Assumed	21,156	25	,846					
	Greenhouse-Geisser	21,156	25,000	,846					
	Huynh-Feldt	21,156	25,000	,846					
	Lower-bound	21,156	25,000	,846					

a. Computed using alpha = ,05

Estimated Marginal Means and Pairwise Comparisons

Grand Mean

Measure: AgentLikeability

		95% Confidence Interval	
Mean	Std. Error	Lower Bound	Upper Bound
3,448	,088	3,267	3,629

Smile

Estimates

Measure: AgentLikeability

			95% Confidence Interval	
Smile	Mean	Std. Error	Lower Bound	Upper Bound
1	3,062	,117	2,820	3,304
2	3,835	,094	3,641	4,029

Pairwise Comparisons

Measure: AgentLikeability

					95% Confidence Interval for Difference[b]	
(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	Lower Bound	Upper Bound
1	2	-,773*	,120	,000	-1,021	-,525
2	1	,773*	,120	,000	,525	1,021

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,623	41,349[a]	1,000	25,000	,000	,623	41,349	1,000
Wilks' lambda	,377	41,349[a]	1,000	25,000	,000	,623	41,349	1,000
Hotelling's trace	1,654	41,349[a]	1,000	25,000	,000	,623	41,349	1,000
Roy's largest root	1,654	41,349[a]	1,000	25,000	,000	,623	41,349	1,000

Each F tests the multivariate effect of Smile. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Gaze

Estimates

Measure: AgentLikeability

Gaze	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	3,303	,110	3,077	3,530
2	3,593	,087	3,413	3,773

Pairwise Comparisons

Measure: AgentLikeability

(I) Gaze	(J) Gaze	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
					Lower Bound	Upper Bound
1	2	-,290[*]	,092	,004	-,479	-,100
2	1	,290[*]	,092	,004	,100	,479

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

	Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
Pillai's trace	,284	9,910[a]	1,000	25,000	,004	,284	9,910	,857
Wilks' lambda	,716	9,910[a]	1,000	25,000	,004	,284	9,910	,857
Hotelling's trace	,396	9,910[a]	1,000	25,000	,004	,284	9,910	,857
Roy's largest root	,396	9,910[a]	1,000	25,000	,004	,284	9,910	,857

Each F tests the multivariate effect of Gaze. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Smile * Subject Extraversion Factor

Estimates

Measure: AgentLikeability

Smile	Subject Extraversion	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	3,270	,209	2,840	3,700
	MEDIUM	2,794	,210	2,361	3,226
	HIGH	3,122	,207	2,695	3,549
2	LOW	3,725	,167	3,380	4,070
	MEDIUM	3,947	,169	3,600	4,294
	HIGH	3,833	,166	3,490	4,176

Pairwise Comparisons

Measure: AgentLikeability

Subject Extraversion	(I) Smile	(J) Smile	Mean Difference (I-J)	Std. Error	Sig.[b]	95% Confidence Interval for Difference[b]	
						Lower Bound	Upper Bound
LOW	1	2	-,455*	,214	,043	-,895	-,015
	2	1	,455*	,214	,043	,015	,895
MEDIUM	1	2	-1,153*	,215	,000	-1,596	-,710
	2	1	1,153*	,215	,000	,710	1,596
HIGH	1	2	-,711*	,212	,003	-1,148	-,274
	2	1	,711*	,212	,003	,274	1,148

Based on estimated marginal means

*. The mean difference is significant at the ,05 level.

b. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Extraversion		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,154	4,538[a]	1,000	25,000	,043	,154	4,538	,535
	Wilks' lambda	,846	4,538[a]	1,000	25,000	,043	,154	4,538	,535
	Hotelling's trace	,182	4,538[a]	1,000	25,000	,043	,154	4,538	,535
	Roy's largest root	,182	4,538[a]	1,000	25,000	,043	,154	4,538	,535
MEDIUM	Pillai's trace	,535	28,731[a]	1,000	25,000	,000	,535	28,731	,999
	Wilks' lambda	,465	28,731[a]	1,000	25,000	,000	,535	28,731	,999
	Hotelling's trace	1,149	28,731[a]	1,000	25,000	,000	,535	28,731	,999
	Roy's largest root	1,149	28,731[a]	1,000	25,000	,000	,535	28,731	,999
HIGH	Pillai's trace	,309	11,205[a]	1,000	25,000	,003	,309	11,205	,896
	Wilks' lambda	,691	11,205[a]	1,000	25,000	,003	,309	11,205	,896
	Hotelling's trace	,448	11,205[a]	1,000	25,000	,003	,309	11,205	,896
	Roy's largest root	,448	11,205[a]	1,000	25,000	,003	,309	11,205	,896

Each F tests the multivariate simple effects of Smile within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

Proxemics * Subject Extraversion Factor

Estimates

Measure: AgentLikeability

Proximity	Subject Extraversion	Mean	Std. Error	95% Confidence Interval	
				Lower Bound	Upper Bound
1	LOW	3,575	,238	3,085	4,066
	MEDIUM	3,089	,240	2,595	3,583
	HIGH	3,646	,237	3,158	4,133
2	LOW	3,420	,182	3,044	3,795
	MEDIUM	3,651	,184	3,273	4,030
	HIGH	3,309	,181	2,936	3,683

Pairwise Comparisons

Measure: AgentLikeability

Subject Extraversion	(I) Proxemics	(J) Proxemics	Mean Difference (I-J)	Std. Error	Sig.[a]	95% Confidence Interval for Difference[a]	
						Lower Bound	Upper Bound
LOW	1	2	,156	,287	,593	-,436	,747
	2	1	-,156	,287	,593	-,747	,436
MEDIUM	1	2	-,562	,289	,063	-1,158	,034
	2	1	,562	,289	,063	-,034	1,158
HIGH	1	2	,336	,286	,250	-,252	,925
	2	1	-,336	,286	,250	-,925	,252

Based on estimated marginal means

a. Adjustment for multiple comparisons: Bonferroni.

Multivariate Tests

Subject Extraversion		Value	F	Hypothesis df	Error df	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[b]
LOW	Pillai's trace	,012	,293[a]	1,000	25,000	,593	,012	,293	,082
	Wilks' lambda	,988	,293[a]	1,000	25,000	,593	,012	,293	,082
	Hotelling's trace	,012	,293[a]	1,000	25,000	,593	,012	,293	,082
	Roy's largest root	,012	,293[a]	1,000	25,000	,593	,012	,293	,082
MEDIUM	Pillai's trace	,131	3,778[a]	1,000	25,000	,063	,131	3,778	,464
	Wilks' lambda	,869	3,778[a]	1,000	25,000	,063	,131	3,778	,464
	Hotelling's trace	,151	3,778[a]	1,000	25,000	,063	,131	3,778	,464
	Roy's largest root	,151	3,778[a]	1,000	25,000	,063	,131	3,778	,464
HIGH	Pillai's trace	,053	1,386[a]	1,000	25,000	,250	,053	1,386	,205
	Wilks' lambda	,947	1,386[a]	1,000	25,000	,250	,053	1,386	,205
	Hotelling's trace	,055	1,386[a]	1,000	25,000	,250	,053	1,386	,205
	Roy's largest root	,055	1,386[a]	1,000	25,000	,250	,053	1,386	,205

Each F tests the multivariate simple effects of Proxemics within each level combination of the other effects shown. These tests are based on the linearly independent pairwise comparisons among the estimated marginal means.

a. Exact statistic

b. Computed using alpha = ,05

A.6 Exploratory Analysis Dataset

The following pages show two tables, one for each trial of the study, including the adjectives reported by subjects after each condition (0-7). Synonyms and adjectives expressing similar impressions are grouped together.

The "Total Count" column shows the total number of hits across all conditions. When a subject gave two or more adjectives falling in the same category, given a single condition, these accounted only for one hit in the respective category.

The hits within individual conditions are listed in the "Condition Counts" columns. Conditions numbers are associated with the levels of our independent variables (smile, gaze and proxemics) in the following manner:

Condition	**Smile**	**Gaze**	**Proxemics**
0	No	Low %	No step
1	No	Low %	Step
2	No	High %	No step
3	No	High %	Step
4	Yes	Low %	No step
5	Yes	Low %	Step
6	Yes	High %	No step
7	Yes	High %	Step

Table A.5: Condition numbers (0-8) and corresponding levels of our independent variables (smile, gaze and proxemics).

First Person Perspective Trial

Adjectives Group	Total Count	Condition Counts							
		0	1	2	3	4	5	6	7
Shy, timid, quiet, withdrawn, introverted, sensitive	26	3	3	6	3	5	2	4	3
Friendly, positive, open	23	2	0	0	1	3	7	4	8
Bored, annoyed, tired	18	6	1	5	3	2	4	2	1
Kind, gentle, polite, thoughtful	20	0	2	3	3	2	3	3	4
Confident, imposing, pushing, achiever, fearless, ambitioned, leader	17	0	7	2	4	0	3	0	1
Aggressive, frightening, stern, unfriendly, contempt, bitter, angry, defensive	14	1	4	3	5	1	0	0	1
Nervous, overanxious, tense, stressed, agitated, impatient, hasty	13	1	3	1	1	2	1	2	3
Happy, Smiley	12	0	0	1	0	2	3	4	2
Composure, professional, business-like, conscientious, task-oriented, accurate, impersonal, discreet, trained	11	1	1	3	3	1	0	3	0
Uninterested, busy, dismissive, careless	10	2	2	2	1	2	3	0	0
Forward, direct, outward directed, open, outgoing, close-talker, extroverted	10	1	2	0	0	2	3	1	2
Helpful, ready to help, interested, people-oriented, eager to meet	9	2	2	0	2	1	2	1	1
Sad, frustrated, depressed, dour	9	2	1	3	1	2	2	0	0
Insecure, confused, unprepared, unconfident, afraid	9	1	1	0	3	0	1	4	0
Neutral, normal, raw, simple	7	3	2	1	1	0	0	3	0
Cold, distant	7	2	0	3	0	2	0	0	2
Strange, weird, unpleasant, creepy	7	2	1	2	0	1	2	0	1
Nice, cute, too sweet	8	0	0	0	1	1	1	3	2
Alive, engaging, responsive, ready	6	1	1	0	0	0	1	1	3
Serious, strict	6	1	2	2	2	0	0	0	0
Calm, slow, relaxed	6	0	0	1	3	1	0	1	0
Welcoming, warm	5	0	1	0	1	0	1	0	2
Unfocused, absent minded, passive, away	3	1	1	0	1	0	0	1	0
Attentive, monitor, supervision, attention, watchful, observant	4	0	1	0	0	2	1	0	0
Funny, humorous	4	0	0	0	0	1	3	0	0
Enthusiastic, excited	3	0	0	0	0	0	1	0	2
Intelligent, brainstorming, concentrated	3	0	0	0	0	1	1	0	1
Agent, scripted, expected, computer-like	3	0	2	1	0	0	0	0	0
Interesting, cool	2	0	0	0	0	0	1	0	1
Looking, seeking, desiring, curious, informant, inquiry, investigative, inquisitive	1	1	0	0	0	0	0	1	0
Security, protection, surveillance, detective, guard, police	2	0	0	1	1	0	0	0	0
Naked	1	0	0	0	0	0	0	0	1
Alone	1	0	0	0	0	0	0	1	0
Unattractive	1	0	0	0	0	1	0	0	0

Third Person Perspective Trial

Adjectives Group	Total Count	Condition Counts							
		0	1	2	3	4	5	6	7
Friendly, open, outgoing, easy-going, positive	30	1	1	1	1	6	6	9	5
Uninterested, busy, indifferent, occupied, careless	23	6	2	6	2	4	2	1	0
Shy, withdrawn, reserved, reticent	20	5	1	2	1	5	1	4	1
Confident, experienced, trustworthy, certain, assured, solid, sure	17	1	3	4	2	3	1	2	1
Aggressive, angry, grumpy, intimidating, hostile, confrontational, challenging, invasive	18	0	6	1	6	0	1	0	4
Helpful, available, interested, involved, willing, forward	16	0	0	2	3	0	4	3	4
Bored, annoyed, tired, uptight, habitual, frustrated, unenthusiastic, unexcited	15	2	1	5	0	3	3	1	0
Absent minded, distracted, daydreamer, preoccupied	14	2	3	1	1	2	2	1	2
Insecure, uncertain, hesitant, unprepared, unsure, untrustworthy	11	3	0	1	1	0	2	2	2
Pushy, bossy, authority, powerful, power, bold, demanding	10	1	3	1	1	0	2	0	2
Composure, professional, business-like, precise, formal, dedicated	10	2	1	4	2	0	0	0	1
Happy, smiley, glad	9	1	0	0	0	3	2	3	0
Excited, active, proactive, enthusiastic	7	0	2	0	0	1	1	0	3
Impatient, eager to meet, nervous	7	2	0	1	3	0	1	0	0
Welcoming, warm, approachable	6	0	1	0	0	0	2	2	1
Fake, unnatural, deliberated, impersonal	6	1	4	0	0	0	1	0	0
Neutral, expressionless, straightforward	6	2	0	2	1	0	1	0	0
Kind, respectful, humble, polite	5	0	1	0	0	0	2	1	1
Calm, relaxed, mellow, patient	5	1	0	2	0	1	0	1	0
Cold, aloof	5	2	0	0	1	0	0	0	2
Nice, mild, pleasing	4	0	0	0	0	0	1	2	1
Present, aware, interactive, awake	3	0	0	0	0	0	0	3	0
Disrespectful, egoistic, pretentious	3	0	1	1	0	0	1	0	0
Afraid, scared, uncomfortable	3	0	0	0	0	1	1	1	0
Old	3	0	0	0	0	1	0	2	0
Sly, sneaky, attentive	2	0	0	0	1	0	1	0	0
Awkward, weird	2	0	0	0	2	0	0	0	0
Unlikeable	1	0	0	1	0	0	0	0	0
Feminine	1	0	1	0	0	0	0	0	0
Interesting	0	0	0	0	0	0	0	0	0

B

Nonverbal Behavior Impact Study Details

B.1 Experimental Room Setting

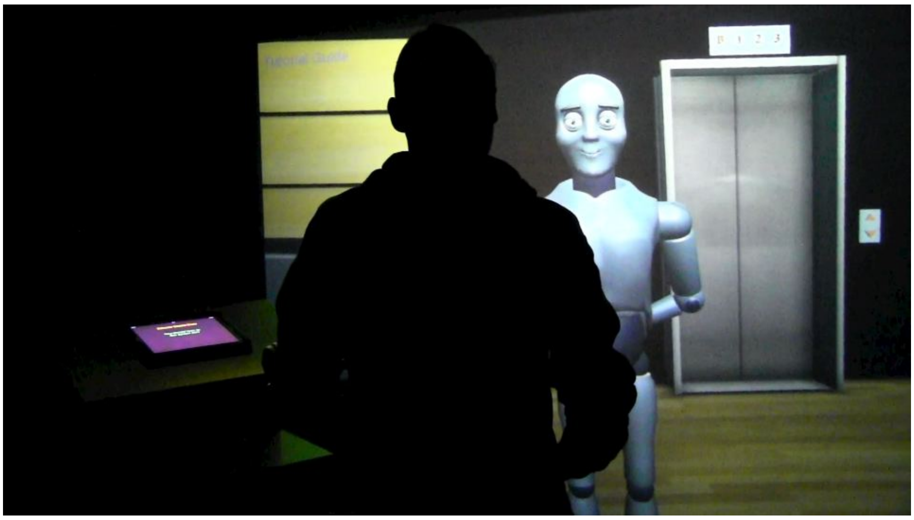

Figure B.1: A picture of the experimental room used for our study showing the user standing still in front of an animated view of the full-size guides.

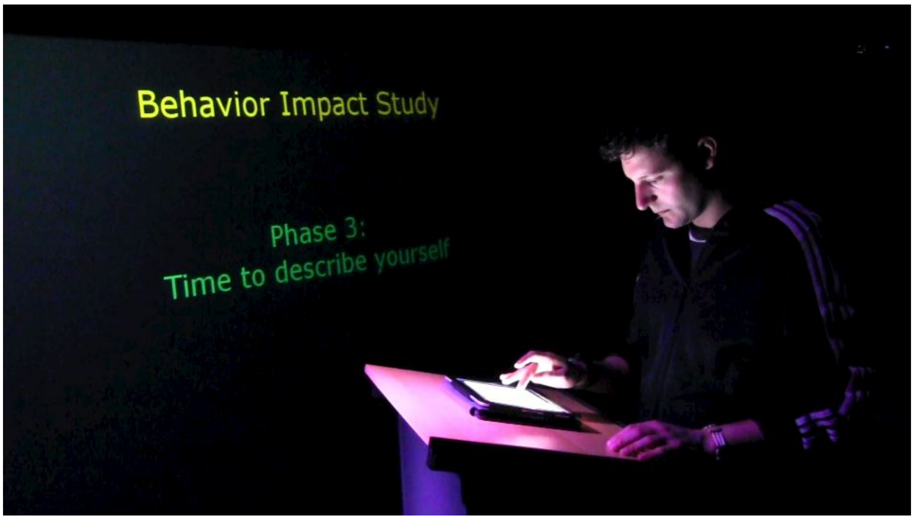

Figure B.2: Another picture of the experimental room showing the user filling the self-report questionnaires after the interaction using the tablet computer on the kiosk.

B.2 Questionnaires

Main Questions

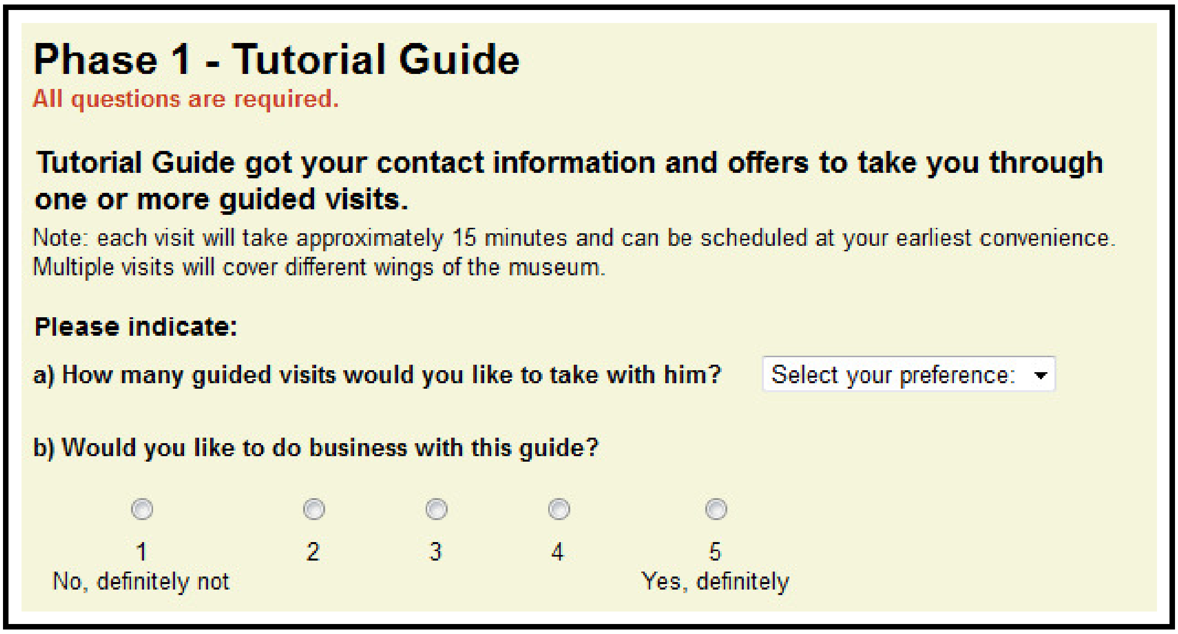

Phase 1 - Tutorial Guide

All questions are required.

Tutorial Guide got your contact information and offers to take you through one or more guided visits.

Note: each visit will take approximately 15 minutes and can be scheduled at your earliest convenience. Multiple visits will cover different wings of the museum.

Please indicate:

a) How many guided visits would you like to take with him? Select your preference: ▾

b) Would you like to do business with this guide?

○	○	○	○	○
1 No, definitely not	2	3	4	5 Yes, definitely

Personality Inventory

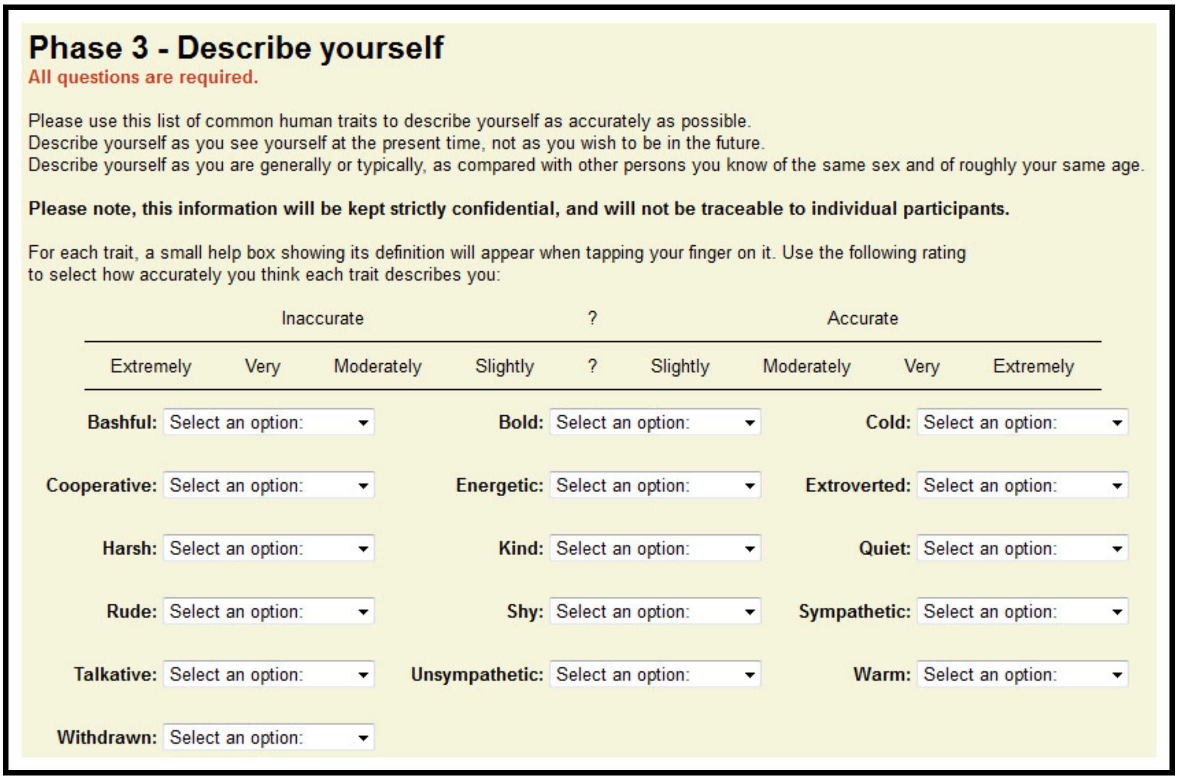

Phase 3 - Describe yourself

All questions are required.

Please use this list of common human traits to describe yourself as accurately as possible.
Describe yourself as you see yourself at the present time, not as you wish to be in the future.
Describe yourself as you are generally or typically, as compared with other persons you know of the same sex and of roughly your same age.

Please note, this information will be kept strictly confidential, and will not be traceable to individual participants.

For each trait, a small help box showing its definition will appear when tapping your finger on it. Use the following rating to select how accurately you think each trait describes you:

Inaccurate				?	Accurate			
Extremely	Very	Moderately	Slightly	?	Slightly	Moderately	Very	Extremely

Bashful: Select an option: ▾ **Bold:** Select an option: ▾ **Cold:** Select an option: ▾

Cooperative: Select an option: ▾ **Energetic:** Select an option: ▾ **Extroverted:** Select an option: ▾

Harsh: Select an option: ▾ **Kind:** Select an option: ▾ **Quiet:** Select an option: ▾

Rude: Select an option: ▾ **Shy:** Select an option: ▾ **Sympathetic:** Select an option: ▾

Talkative: Select an option: ▾ **Unsympathetic:** Select an option: ▾ **Warm:** Select an option: ▾

Withdrawn: Select an option: ▾

Demographics

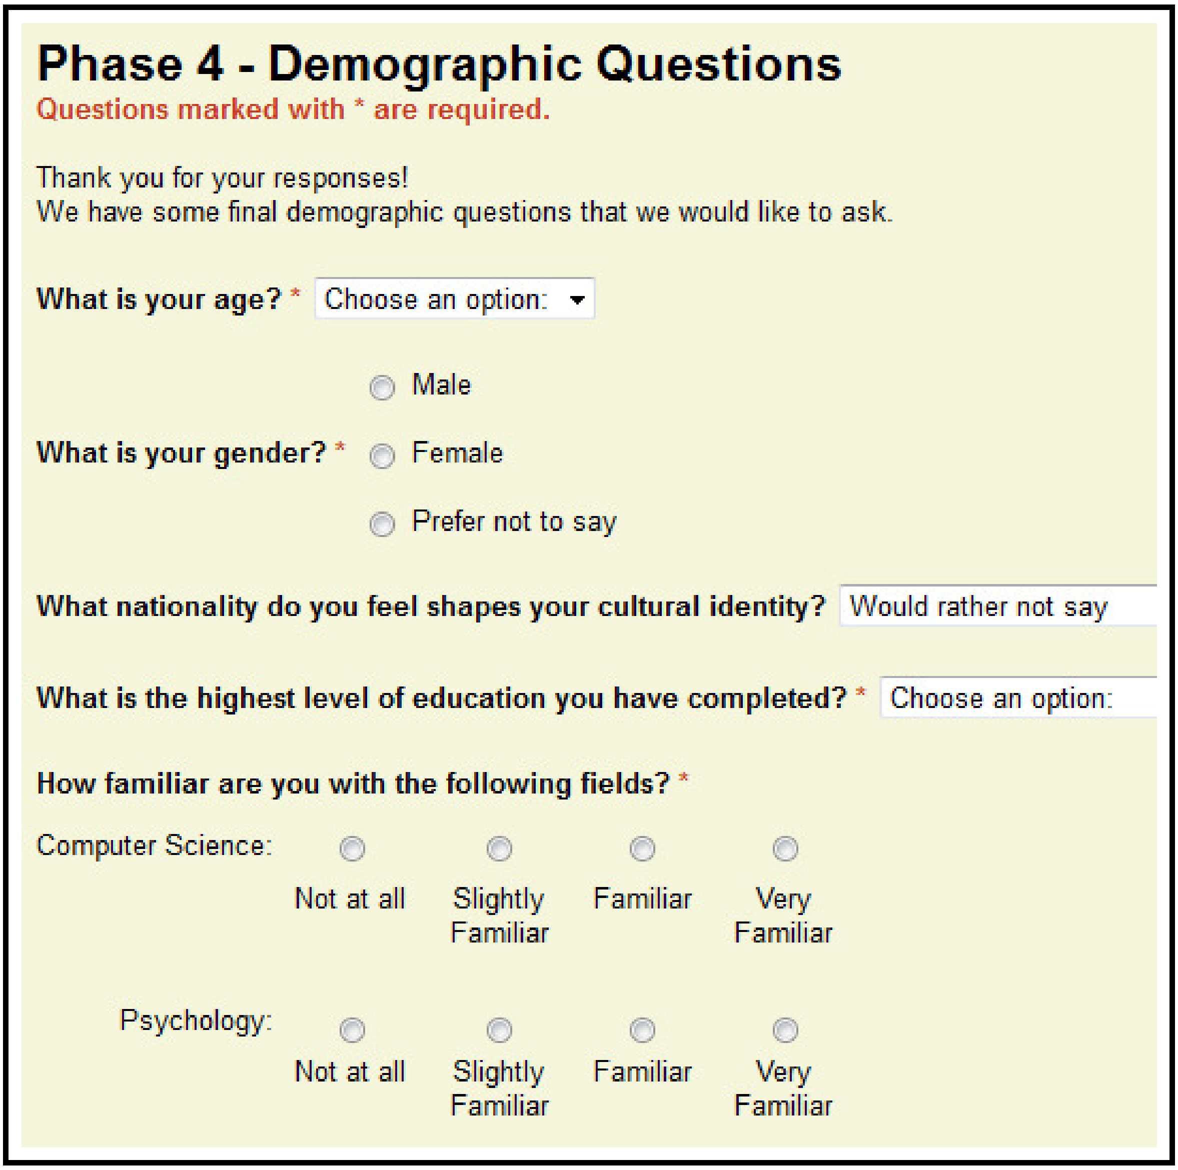

Phase 4 - Demographic Questions

Questions marked with * are required.

Thank you for your responses!
We have some final demographic questions that we would like to ask.

What is your age? * Choose an option:

What is your gender? *
○ Male
○ Female
○ Prefer not to say

What nationality do you feel shapes your cultural identity? Would rather not say

What is the highest level of education you have completed? * Choose an option:

How familiar are you with the following fields? *

Computer Science:	○ Not at all	○ Slightly Familiar	○ Familiar	○ Very Familiar
Psychology:	○ Not at all	○ Slightly Familiar	○ Familiar	○ Very Familiar

B.3 Demographics

Subjects Age

Range	Frequency
< 18	-
18-20	-
21-30	13
31-40	8
41-50	2
51-60	1
> 61	-
Total	24

Table B.1: Subjects age frequencies in range intervals (shown in years).

Subjects Cultural Identity

Country	Frequency
Canada	1
Denmark	1
Iceland	15
Italy	1
Romania	1
Spain	2
Sweden	1
United Kingdom	1
Not said	1
Total	24

Table B.2: Subjects cultural identity frequencies by country.

Subjects Level of Education

Education Level	**Frequency**
Doctorate level	3
Master level	8
Undergraduate level	6
High school	6
Elementary school	-
Less than elementary school	-
Other	1
Total	32

Table B.3: Subjects level of education frequencies.

B.4 Documents

Consent Declaration Form

Behavior Impact Study

CONSENT DECLARATION
May/June 2012

Principal Investigator:
Angelo Cafaro, CADIA, School of Computer Science, Reykjavik University

In this study, titled "Behavior Impact Study", I will be interacting with some graphical agents in the main reception of a virtual museum. These will be shown on a 2.70x2.00 (m) back-projected screen and I will use an iPad for controlling the interaction and answering questions after meeting each of them. I fully understand that my participation is voluntary and that I am free to withdraw my consent and to discontinue participation at *any time* without prejudice to myself. The experimental procedure has been explained to me and the investigator, Angelo Cafaro, has offered to answer any inquiries concerning this procedure.

By signing this document, I agree to come back in these facilities and be guided in one or more virtual tours of the museum according to the particular guide that will be assigned to me and the preference (number of visits) that I expressed for him. The guided tours (if any) will be scheduled over a period of two consecutive months at my earliest convenience. The start date of this period and all appointments will be scheduled in concordance with the investigator. Every visit will require approximately 15 minutes and different area of the museum will be shown if more than one visit is scheduled.

I have the option to leave my e-mail address, by doing so I accept to participate to a lottery involving all subjects participating in this study and I allow the Investigator to contact me by e-mail in case I'll be drawn. The prize for this lottery are two tickets for a movie theater in Reykjavik and participation is free of any charge.

I understand that I can contact the Director of Research Services at Reykjavik University, Kristján Kristjánsson (kkru.is), if I believe I have been treated unfairly as a subject and/or I believe that the research team has breached the RU Code of Ethics.

I have read and understood the above, as well as the experiment instructions, and agree to participate in this research effort.

FULL NAME OF SUBJECT

E-MAIL (OPTIONAL)

SIGNATURE and DATE

Instructions

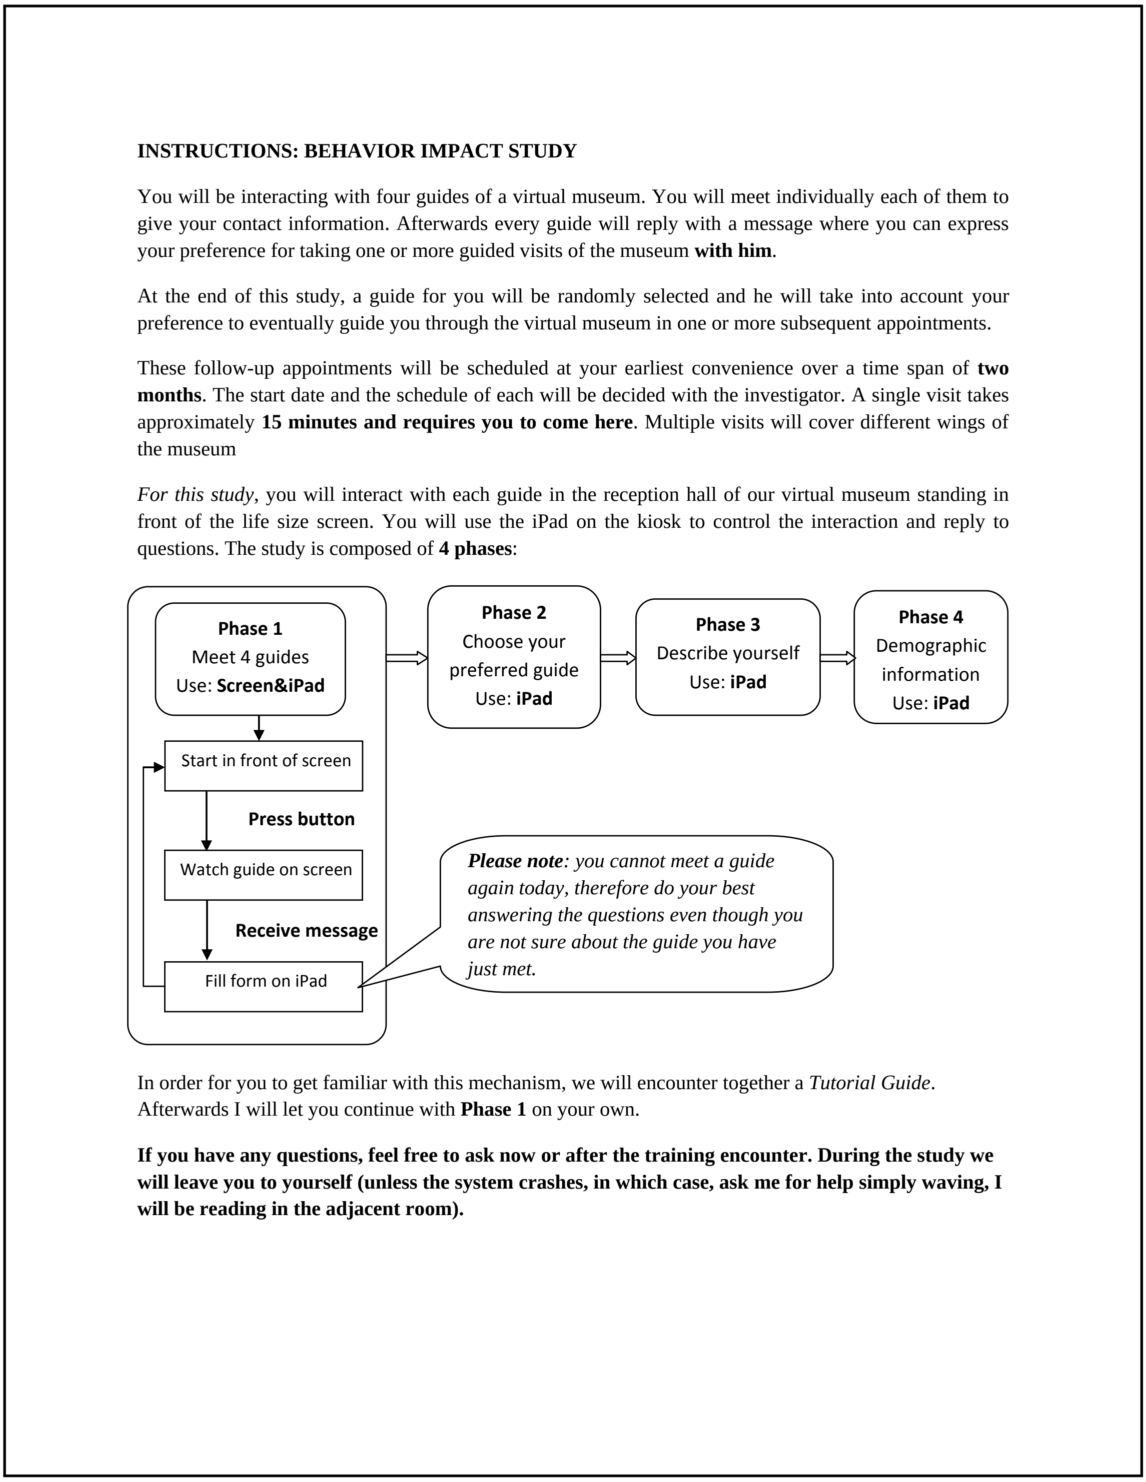

INSTRUCTIONS: BEHAVIOR IMPACT STUDY

You will be interacting with four guides of a virtual museum. You will meet individually each of them to give your contact information. Afterwards every guide will reply with a message where you can express your preference for taking one or more guided visits of the museum **with him**.

At the end of this study, a guide for you will be randomly selected and he will take into account your preference to eventually guide you through the virtual museum in one or more subsequent appointments.

These follow-up appointments will be scheduled at your earliest convenience over a time span of **two months**. The start date and the schedule of each will be decided with the investigator. A single visit takes approximately **15 minutes and requires you to come here**. Multiple visits will cover different wings of the museum

For this study, you will interact with each guide in the reception hall of our virtual museum standing in front of the life size screen. You will use the iPad on the kiosk to control the interaction and reply to questions. The study is composed of **4 phases**:

***Please note**: you cannot meet a guide again today, therefore do your best answering the questions even though you are not sure about the guide you have just met.*

In order for you to get familiar with this mechanism, we will encounter together a *Tutorial Guide*. Afterwards I will let you continue with **Phase 1** on your own.

If you have any questions, feel free to ask now or after the training encounter. During the study we will leave you to yourself (unless the system crashes, in which case, ask me for help simply waving, I will be reading in the adjacent room).

Debriefing

Behavior Interpretation Study

DEBRIEFING
May/June 2012

Principal Investigator:
Angelo Cafaro, CADIA, School of Computer Science, Reykjavik University

You have just participated in a study intended to help understanding how your first impressions of a virtual character (the guides you have met) impact your decision regarding how often you choose to spend time with him later. We have also studied possible matches between your own personality and the choices you have made.

Each guide you have been approaching was programmed to manage his impressions of extraversion (high vs. low) and friendliness (hostile vs. friendly). In particular, the combination of these two characteristics was obtained by exhibiting only nonverbal immediacy cues of Smile, Gaze and Proximity during the approaches.

The questions you replied after being contacted by each guide will be used (a) to study your level of commitment to interact again with that particular agent and (b) in general if you like to do business with him at all. In the end, the personality inventory you have filled in, will be used to see if there exists any kind of correlation between the previous ratings and your own personality.

Your personal information will be kept strictly confidential, will not be traceable to individual participants and will not be sold, reused, rented, loaned or otherwise disclosed. Any information you give us will not be used in ways that you have not consented to.

The visits offered are fictitious, we will *not* schedule any guided visit nor randomly assign any guide to you. This scenario was introduced to give importance to your choices adding a true time commitment to them. If you, for any reason, are not comfortable with the way the experiment was conducted, feel free to withdraw your consent declaration and cancel your participation.

We kindly ask you to keep secret this short explanation of the study purposes, since we will be running this study on other persons you might know or discuss the study with.

B.5 Summary of Means and ANOVA Tables

Likelihood of Encounters Measure

Tests of Within-Subjects Effects

Measure: Likelihood of Encounters

		Type III Sum of Squares	df	Mean Square	F	Sig.	Partial Eta Squared	Noncent. Parameter	Observed Power[a]
Guide Extraversion	Sphericity Assumed	,407	1	,407	,390	,539	,018	,390	,092
	Greenhouse-Geisser	,407	1,000	,407	,390	,539	,018	,390	,092
	Huynh-Feldt	,407	1,000	,407	,390	,539	,018	,390	,092
	Lower-bound	,407	1,000	,407	,390	,539	,018	,390	,092
Guide Extraversion * Subject Extraversion	Sphericity Assumed	,253	1	,253	,243	,627	,011	,243	,076
	Greenhouse-Geisser	,253	1,000	,253	,243	,627	,011	,243	,076
	Huynh-Feldt	,253	1,000	,253	,243	,627	,011	,243	,076
	Lower-bound	,253	1,000	,253	,243	,627	,011	,243	,076
Guide Extraversion * Subject Agreeableness	Sphericity Assumed	,110	1	,110	,105	,749	,005	,105	,061
	Greenhouse-Geisser	,110	1,000	,110	,105	,749	,005	,105	,061
	Huynh-Feldt	,110	1,000	,110	,105	,749	,005	,105	,061
	Lower-bound	,110	1,000	,110	,105	,749	,005	,105	,061
Error(Guide Extraversion)	Sphericity Assumed	21,925	21	1,044					
	Greenhouse-Geisser	21,925	21,000	1,044					
	Huynh-Feldt	21,925	21,000	1,044					
	Lower-bound	21,925	21,000	1,044					
Guide Friendliness	Sphericity Assumed	24,204	1	24,204	21,909	,000	,511	21,909	,994
	Greenhouse-Geisser	24,204	1,000	24,204	21,909	,000	,511	21,909	,994
	Huynh-Feldt	24,204	1,000	24,204	21,909	,000	,511	21,909	,994
	Lower-bound	24,204	1,000	24,204	21,909	,000	,511	21,909	,994
Guide Friendliness * Subject Extraversion	Sphericity Assumed	1,014	1	1,014	,917	,349	,042	,917	,150
	Greenhouse-Geisser	1,014	1,000	1,014	,917	,349	,042	,917	,150
	Huynh-Feldt	1,014	1,000	1,014	,917	,349	,042	,917	,150
	Lower-bound	1,014	1,000	1,014	,917	,349	,042	,917	,150
Guide Friendliness * Subject Agreeableness	Sphericity Assumed	,290	1	,290	,262	,614	,012	,262	,078
	Greenhouse-Geisser	,290	1,000	,290	,262	,614	,012	,262	,078
	Huynh-Feldt	,290	1,000	,290	,262	,614	,012	,262	,078
	Lower-bound	,290	1,000	,290	,262	,614	,012	,262	,078
Error(Guide Friendliness)	Sphericity Assumed	23,200	21	1,105					
	Greenhouse-Geisser	23,200	21,000	1,105					
	Huynh-Feldt	23,200	21,000	1,105					
	Lower-bound	23,200	21,000	1,105					
Guide Extraversion * Guide Friendliness	Sphericity Assumed	,461	1	,461	,633	,435	,029	,633	,118
	Greenhouse-Geisser	,461	1,000	,461	,633	,435	,029	,633	,118
	Huynh-Feldt	,461	1,000	,461	,633	,435	,029	,633	,118
	Lower-bound	,461	1,000	,461	,633	,435	,029	,633	,118
GuideExtraversion * GuideFriendliness * SExtraversionFactor	Sphericity Assumed	,599	1	,599	,821	,375	,038	,821	,139
	Greenhouse-Geisser	,599	1,000	,599	,821	,375	,038	,821	,139
	Huynh-Feldt	,599	1,000	,599	,821	,375	,038	,821	,139
	Lower-bound	,599	1,000	,599	,821	,375	,038	,821	,139
Guide Extraversion * Guide Friendliness * Subject Agreeableness	Sphericity Assumed	,183	1	,183	,251	,622	,012	,251	,077
	Greenhouse-Geisser	,183	1,000	,183	,251	,622	,012	,251	,077
	Huynh-Feldt	,183	1,000	,183	,251	,622	,012	,251	,077
	Lower-bound	,183	1,000	,183	,251	,622	,012	,251	,077
Error(Guide Extraversion * Guide Friendliness)	Sphericity Assumed	15,314	21	,729					
	Greenhouse-Geisser	15,314	21,000	,729					
	Huynh-Feldt	15,314	21,000	,729					
	Lower-bound	15,314	21,000	,729					

a. Computed using alpha = .05

Estimated Marginal Means

Grand Mean

Measure: Likelihood of Encounters

Mean	Std. Error	95% Confidence Interval	
		Lower Bound	Upper Bound
2,943	,089	2,758	3,128

Guide Extraversion

Measure: Likelihood of Encounters

Guide Extraversion	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	3,008	,140	2,717	3,300
2	2,877	,135	2,596	3,158

Guide Friendliness

Measure: Likelihood of Encounters

Guide Friendliness	Mean	Std. Error	95% Confidence Interval	
			Lower Bound	Upper Bound
1	2,436	,137	2,152	2,721
2	3,449	,143	3,151	3,747

Number of Visits Measure

This measure was on an ordinal scale (0-10). A parametric analysis for this measure would be the repeated measures ANOVA. The Aligned Rank Transform (ART) procedure was devised to cope with this issue in case of non-parametric data [Wobbrock et al., 2011]. With ART data are aligned, ranked and then analyzed with the appropriate parametric procedure. For each main effect or interaction, the response variable Y (***Number of Visits*** in our case) was "aligned", a process that "strips" from Y all effects but the one of interest. This aligned response is called $Y_{aligned}$. The aligned responses were then assigned ranks and averaged in the case of ties. Thes new responses were called Y_{art}. Then a full factorial ANOVA was run on the Y_{art} responses, but only the effect for which Y was aligned it was examined. Thus, for each main effect or interaction, a new aligned column ($Y_{aligned}$) and a new ranked column (Y_{art}) were necessary. In our study, for example, with two factors (***Guide Extraversion*** and ***Guide Friendliness***) and their interaction, we needed six additional columns, three aligned and three ranked, where each set of three comprised each of two factors and their interaction. In general, for N factors, 2^{N-1} aligned columns and 2^{N-1} ranked columns are needed. It is possible to study further interaction effects with between subject factors (in our case ***Subject Extraversion*** and ***Agreeablness*** traits) by including them one at time in a mixed design ANOVA as explained in Section 4.3.5. The full ANOVA tables after aligning and ranking the data is omitted from this appendix but it is available upon request by contacting the author.

C

Managing First Impressions in a Public Setting Study Details

C.1 Tinker's Exhibit at the Museum of Science

Figure C.1: Tinker is exhibited in the main entrance of the Computer Place area at the Boston Museum of Science in USA, MA. The exhibit became operational in September, 2007.

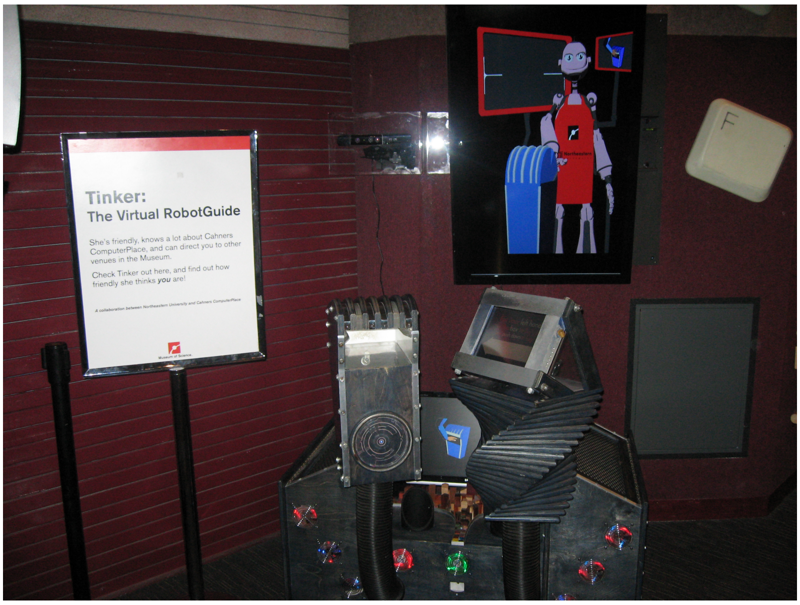

Figure C.2: This picture shows the setup of the exhibit. The agent is projected onto a 46" LCD screen, it can detect approaching visitors' proximity in real-time with the Microsoft's kinect mounted aside the screen, the biometric hand reader (on the left under the screen) can re-identify visitors it has already talked to, and the multiple choice touch screen input (on the right) allows visitors to exchange dialogue turns with the agent.

C.2 Manipulation Check

We ran an informal manipulation check following a procedure similar to the one described in section A.1 for the nonverbal behavior interpretation study. Though we only had one smiling animation available, there were two options for the gaze behavior the we intended to test in order to chose the best one for our study. In one version Tinker was gazing away from the user by looking at its right with the eyes (GAZE AWAY RIGHT), whereas in the other version the gaze away was done by looking down with the eyes only at the virtual hand reader in its 3D environment (GAZE AWAY DOWN). In summary, our goals for this manipulation check were:

1. Validate the nonverbal cues exhibited by Tinker by making sure that subjects correctly perceived the difference between the two levels of each one (i.e. smiling vs. not smiling, more gazing at user vs. fewer gazing);
2. Choose the gaze animation version;
3. Ensure that the agent's graphics were visually clear.

Setting and Procedure

We had 10 colleagues in our lab (3 females and 7 males) that approached Tinker in a setup similar to the original one at the museum, but in our lab facilities. This setup included only the screen to display the agent and the kinect to sense their proximity.

Subjects were randomly assigned to an initial GAZE AWAY VERSION of Tinker. The two manipulations (i.e. smile and gaze nonverbal cues) were exhibited in isolation from each other. Since there were two levels for each one, subjects performed a succession of two approaches for each manipulation (in a fully randomized order), and then replied to a question comparing the pair, as described in the following section. After completing all the pairs for a given GAZE AWAY VERSION, they repeated the same procedure for the other one.

Questions

We used paper and pencil to administer the questionnaires by gradually disclosing the full question set using separate sheets to avoid priming effects. The questions administered are listed below. For questions 1-2 there was a follow-up question in case subjects were able to observe a difference (i.e. replying *"Yes"*) between the two approaches (possible answers for these cases are indicated in square brackets). The last two questions were shown after completing all the pairs in a GAZE AWAY VERSION. They were used to address graphics related or general issues (question 3), and to gather general comments and feedbacks (question 4).

1. **Did the agent smile towards you during the approaches?**
    - No
    - Yes. If so, which one smiled at you? [First | Second | Both]
2. **Did the agent look at you more in one of the approaches?**
    - No
    - Yes. If so, in which one it looked at you more? [First | Second]
3. **In general, did you see how the agent reacted towards you?**
    - Yes
    - No, the graphics were visually unclear
    - No, because (use comments box)
4. **Any comments?**

Summary of Answers

A summary of subjects' answers grouped by GAZE AWAY VERSION is provided in table C.1. The **"No diff."** column represents total number of cases where subjects weren't able to observe any difference. Then we counted the total number of **"Match"** cases. In other words, when in the first approach of a given pair for a manipulation, for example the smile, Tinker was smiling and they correctly marked *"First"*, we counted this as a match. Vice-versa, the **"Mismatch"** column includes cases where subjects attributed the behavior in question to the wrong approach[1]. The total counts in the **"General Question"** raw indicates whether Tinker's nonverbal reactions were visible in general (**Reaction Visible**), if there were problems with the agent's graphics (**Graphics Problem**) or other problems later reported in the comments box (**Other Problem**).

Gaze Away	**Manipulation**	**No diff.**	**Match**	**Mismatch**
RIGHT	Smile	0	10	0
	Gaze	1	6	3
	General Question	**Reaction Visible**	**Graphics Problem**	**Other Problems**
		10	0	0
Gaze Away	**Manipulation**	**No diff.**	**Match**	**Mismatch**
DOWN	Smile	0	8	2
	Gaze	1	9	0
	General Question	**Reaction Visible**	**Graphics Problem**	**Other Problems**
		10	0	0

Table C.1: Summary of answers provided by subjects in the manipulation check grouped by GAZE AWAY VERSION (**Gaze Away**). For the two manipulations, the columns indicate total counts of cases where subjects: didn't see any difference (**No diff.**), correctly matched their answer with the levels observed in succession (**Match**) and, vice-versa, mis-attributed a behavior in a pair of approaches (**Mismatch**). The **General Question** raw reports total counts for question 3.

From these responses emerged that, in terms of gaze behavior, the version of Tinker looking down the virtual hand reader when gazing away from the subject was perceived slightly better than the one gazing away at its right. However, we had some concerns about the smiling animation. Although there seemed to be no problems in both versions tested, the particular smiling animation that Tinker was capable of exhibit in the GAZE AWAY DOWN version was a bit mechanical due to some Tinker's legacy animations constraints. Whereas it was more natural in the GAZE AWAY RIGHT version. We also had some subjects reporting that *"detecting the smile was a little hard in gaze away down version"* and *"the gaze away right*

---

[1] The count of mismatches for question 1 includes also cases where subject replied with "Both".

is better". So, given that smile had great impact in both our previous studies in terms of managing impressions of friendliness, we gave priority to this nonverbal cue and we adopted the GAZE AWAY RIGHT version of Tinker for the main study at the museum.

In conclusion, subjects were able to observe differences in the manipulations observed during the approaches and there weren't graphics related issues with Tinker. There weren't other major issues from their feedbacks.

C.3 Web Survey Demographics

Subjects Age

Range	**Frequency**
< 18	-
18-20	10
21-30	62
31-40	34
41-50	17
51-60	3
> 61	-
Total	126

Table C.2: Subjects age frequencies in range intervals (shown in years).

Subjects Cultural Identity

Country	**Frequency**
Australia	1
France	1
Iceland	110
Italy	3
Sweden	1
United States	1
Not said	9
Total	126

Table C.3: Subjects cultural identity frequencies by country.

Subjects Level of Education

Education Level	**Frequency**
Doctorate level	1
Master level	25
Undergraduate level	44
High school	52
Elementary school	1
Less than elementary school	-
Other	3
Total	126

Table C.4: Subjects level of education frequencies.

C.4 Web Survey Exploratory Analysis Dataset

The following table shows the adjectives reported by subjects after watching each video of Tinker showing the HOSTILE and FRIENDLY version respectively. Synonyms and adjectives expressing similar impressions are grouped together.

The "Total Count" column shows the total number of hits across the two conditions. When a subject gave two or more adjectives falling in the same category, given a single condition, these accounted only for one hit in the respective category. The hits within a given condition are listed in the "Condition Counts" columns.

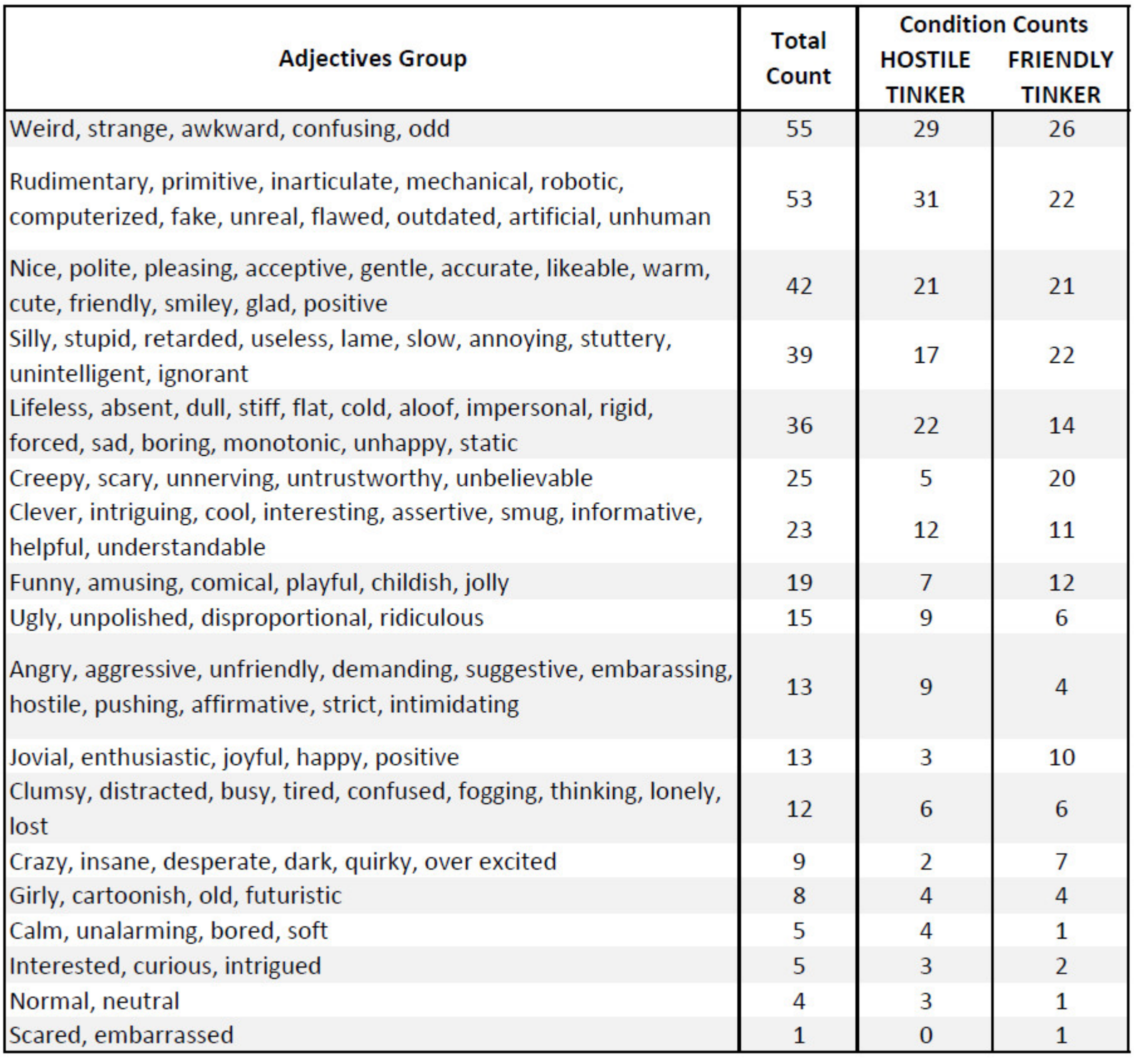

Adjectives Group	Total Count	Condition Counts	
		HOSTILE TINKER	FRIENDLY TINKER
Weird, strange, awkward, confusing, odd	55	29	26
Rudimentary, primitive, inarticulate, mechanical, robotic, computerized, fake, unreal, flawed, outdated, artificial, unhuman	53	31	22
Nice, polite, pleasing, acceptive, gentle, accurate, likeable, warm, cute, friendly, smiley, glad, positive	42	21	21
Silly, stupid, retarded, useless, lame, slow, annoying, stuttery, unintelligent, ignorant	39	17	22
Lifeless, absent, dull, stiff, flat, cold, aloof, impersonal, rigid, forced, sad, boring, monotonic, unhappy, static	36	22	14
Creepy, scary, unnerving, untrustworthy, unbelievable	25	5	20
Clever, intriguing, cool, interesting, assertive, smug, informative, helpful, understandable	23	12	11
Funny, amusing, comical, playful, childish, jolly	19	7	12
Ugly, unpolished, disproportional, ridiculous	15	9	6
Angry, aggressive, unfriendly, demanding, suggestive, embarassing, hostile, pushing, affirmative, strict, intimidating	13	9	4
Jovial, enthusiastic, joyful, happy, positive	13	3	10
Clumsy, distracted, busy, tired, confused, fogging, thinking, lonely, lost	12	6	6
Crazy, insane, desperate, dark, quirky, over excited	9	2	7
Girly, cartoonish, old, futuristic	8	4	4
Calm, unalarming, bored, soft	5	4	1
Interested, curious, intrigued	5	3	2
Normal, neutral	4	3	1
Scared, embarrassed	1	0	1

D

FML Standard 1.0

This Appendix contains the current FML specification proposal. The specification describes a document divided in two main sections. The first section is named **declarations**. It is used to identify the participants referred by the functions in the body section and to specify contextual information such as the floors configuration. The elements and attributes that are supported in this section of the document are illustrated in Section D.1.

The second section is named **body**. It is used to specify the functions that each participant needs to accomplish and it is divided into three separate tracks named: *interactional*, *performative* and *mental-state*. Each track hosts different elements that describe functions meaningful for the track they belong to. The elements in each track are organized in FML Chunks and timed with relative temporal constraints. These common elements definitions (FML chunk and temporal constraints) and specific elements appearing in each track are illustrated in section D.2.

The root element of an FML document is the **`<saiba-act>`** tag with the following namespace: "**http://cadia.ru.is/FMLSpecification**". This element allows the inclusion of a **`<declarations>`** and a **`<body>`** element.

D.1 FML Document Declaration

element **DECLARATIONS**

Description	Stores contextual information.				
Attributes	Name	Type	Use	Default	Description
	-	-	-	-	-
Elements	Name	Occurs	Description		
	IDENTIKITS	1..1	Contains an `<identikit>` for each participant.		
	FLOORS	1..1	Contains a `<floor>` for each active floor that is described by the FML document.		

D.1.1 Identikits Information

element **IDENTIKITS**

Description	Stores participants' identikits.				
Attributes	Name	Type	Use	Default	Description
	-	-	-	-	-
Elements	Name	Occurs	Description		
	IDENTIKIT	1..*	An `<identikit>` for each participant.		

element **IDENTIKIT**

<table>
<tr><td>Description</td><td colspan="5">A participant’s identikit.</td></tr>
<tr><td rowspan="4">Attributes</td><td>Name</td><td>Type</td><td>Use</td><td>Default</td><td>Description</td></tr>
<tr><td>ID</td><td><i>ID</i></td><td>required</td><td>-</td><td>Unique identifier.</td></tr>
<tr><td>NAME</td><td><i>string</i></td><td>required</td><td>-</td><td>Participant’s name.</td></tr>
<tr><td>GENDER</td><td><i>enum</i></td><td>optional</td><td>-</td><td>Participant’s gender. Values [<i>male</i> | <i>female</i>].</td></tr>
<tr><td rowspan="3">Elements</td><td>Name</td><td>Occurs</td><td colspan="3">Description</td></tr>
<tr><td>PERSONALITY</td><td>0..1</td><td colspan="3">Participant’s personality.</td></tr>
<tr><td>RELATIONSHIPS</td><td>0..1</td><td colspan="3">Contains information about the participant’s relationship levels with others.</td></tr>
</table>

element **PERSONALITY**

<table>
<tr><td>Description</td><td colspan="5">The participant’s personality traits.</td></tr>
<tr><td rowspan="6">Attributes</td><td>Name</td><td>Type</td><td>Use</td><td>Default</td><td>Description</td></tr>
<tr><td>EXTRAVERSION</td><td><i>enum</i></td><td>optional</td><td>-</td><td>Values: [<i>LOW</i> | <i>NEUTRAL</i> | <i>HIGH</i>].</td></tr>
<tr><td>AGREEABLENESS</td><td><i>enum</i></td><td>optional</td><td>-</td><td>″</td></tr>
<tr><td>NEUROTICISM</td><td><i>enum</i></td><td>optional</td><td>-</td><td>″</td></tr>
<tr><td>CONSCIENTIOUSNESS</td><td><i>enum</i></td><td>optional</td><td>-</td><td>″</td></tr>
<tr><td>OPENNESS</td><td><i>enum</i></td><td>optional</td><td>-</td><td>″</td></tr>
<tr><td rowspan="2">Elements</td><td>Name</td><td>Occurs</td><td colspan="3">Description</td></tr>
<tr><td>-</td><td>-</td><td colspan="3">-</td></tr>
</table>

element **RELATIONSHIPS**

Description	The participant's relationship levels with other participants.				
Attributes	Name	Type	Use	Default	Description
	-	-	-	-	-
Elements	Name	Occurs	Description		
	RELATIONSHIP	1..*	A series of relationship elements.		

element **RELATIONSHIP**

Description	A participant's relationship level.				
Attributes	Name	Type	Use	Default	Description
	LEVEL	*enum*	required	-	Relationship level. Values [*STRANGER* | *ACQUAINTANCE* | *FRIEND*].
	WITH	*IDREF*	required	-	Reference to other participant.
Elements	Name	Occurs	Description		
	-	-	-		

D.1.2 Floors Information

element **FLOORS**

Description	Stores floors' information.
Attributes	Name / Type / Use / Default / Description - / - / - / - / -
Elements	Name / Occurs / Description **FLOOR** / 1..* / A `<floor>` for each active floor described in the document.

element **FLOOR**

	Name	Type	Use	Default	Description
Description	A floor's configuration.				
Attributes	**FLOOR-ID**	*ID*	required	-	Unique identifier.
	FLOOR-CFG	*enum*	required	-	Floor's configuration. Values [*individual* | *unicast* | *broadcast* | *multicast*].
Elements	**Name**	**Occurs**	**Description**		
	PARTICIPANT	1..*	Participant's contextual information.		

element **PARTICIPANT**

	Name	Type	Use	Default	Description
Description	Participant's contextual information.				
Attributes	**IDENTIKIT-REF**	*IDREF*	required	-	Reference to identikit.
	ENTITY	*enum*	required	-	Participant's entity. Values [*individual* | *group*].
	ROLE	*enum*	required	-	Participant's role. Values [*speaker* | *addressed-hearer* | *unaddressed-hearer* | *eavesdropper* | *overhearer*].
Elements	**Name**	**Occurs**	**Description**		
	ATTITUDE	0..*	Participant's attitudes towards others.		

element **ATTITUDE**

Description	A participant's attitude towards another participant.				
Attributes	Name	Type	Use	Default	Description
	AFFILIATION	*enum*	required	-	Affiliation level. Values: [*HOSTILE* | *NEUTRAL* | *FRIENDLY*].
	STATUS	*enum*	required	-	Status level. Values: [*DOMINANT* | *NEUTRAL* | *SUBMISSIVE*].
	TOWARDS	*IDREF*	required	-	Reference to another participant.
Elements	Name	Occurs	Description		
	-	-	-		

D.2 FML Document Body

element **BODY**

Description	The body of an FML document divided in three tracks.				
Attributes	Name	Type	Use	Default	Description
	-	-	-	-	-
Elements	Name	Occurs	Description		
	INTERACTIONAL	1..1	Contains interactional functions.		
	PERFORMATIVE	1..1	Contains performative functions.		
	MENTAL-STATE	1..1	Contains mental state functions.		

D.2.1 Common Body Elements and Attributes

The functions featured in each track of the body are organized in FML Chunks and timed with relative temporal constraints. Therefore, each track (*interactional*, *performative* or *mental-state*) can include *zero* or *more* occurrences of `<fml-chunk>` elements. Prior to describing all the FML elements relative to the functions that each track can host, we illustrate the specifications for **FML chunk** and **temporal constraint** elements.

element **FML-CHUNK**

Description	FML chunk elements included in the three body tracks.				
Attributes	Name	Type	Use	Default	Description
	ACTID	*ID*	required	-	Unique identifier.
	PARTICIPANT-REF	*IDREF*	required	-	Reference to participant's identikit.
Elements	Name	Occurs	Description		
	TIMING	0..1	Chunk's temporal constraint.		
	function element	1..*	Track specific functions.		

element **TIMING**

Description	Temporal constraint descriptor.				
Attributes	Name	Type	Use	Default	Description
	PRIMITIVE	*enum*	required	-	The temporal constraint. Values [*immediately* | *must_end_before* | *execute_anytime_during* | *start_immediately_after* | *start_sometime_after* | *start_together*].
	ACTREF	*IDREF*	optional	-	Reference to another FML chunk when required by the primitive.
Elements	Name	Occurs	Description		
	-	-	-		

The following sections describe, separately for each track, the elements representing the functions that the tracks can host. These elements have two attributes in common that we describe here as follows. We will omit these two attributes from the descriptions of the elements in the remainder of this appendix:

- *floorID:* A reference to the floor in which the function described by the element is meant to be accomplished;
- *id:* A unique identifier associated with the function.

D.2.2 Interactional Track Elements

element **INITIATION**

<table>
<tr><td>Description</td><td colspan="5">Initiation of interaction category.</td></tr>
<tr><td rowspan="3">Attributes</td><td>Name</td><td>Type</td><td>Use</td><td>Default</td><td>Description</td></tr>
<tr><td>TYPE</td><td><i>enum</i></td><td>required</td><td>-</td><td>The specific type.
Values [<i>react</i> | <i>recognize</i> | <i>salute-distant</i> | <i>approach-react</i> | <i>salute-close</i> | <i>initiate</i>].</td></tr>
<tr><td>ADDRESSEE</td><td><i>IDREF</i></td><td>required</td><td>-</td><td>Reference to the addressee of this function.</td></tr>
<tr><td rowspan="2">Elements</td><td>Name</td><td>Occurs</td><td colspan="3">Description</td></tr>
<tr><td>-</td><td>-</td><td colspan="3">-</td></tr>
</table>

element **CLOSING**

<table>
<tr><td>Description</td><td colspan="5">Closing of interaction category.</td></tr>
<tr><td rowspan="3">Attributes</td><td>Name</td><td>Type</td><td>Use</td><td>Default</td><td>Description</td></tr>
<tr><td>TYPE</td><td><i>enum</i></td><td>required</td><td>-</td><td>The specific type.
Values [<i>break-away</i> | <i>farewell</i>].</td></tr>
<tr><td>ADDRESSEE</td><td><i>IDREF</i></td><td>required</td><td>-</td><td>Reference to the addressee of this function.</td></tr>
<tr><td rowspan="2">Elements</td><td>Name</td><td>Occurs</td><td colspan="3">Description</td></tr>
<tr><td>-</td><td>-</td><td colspan="3">-</td></tr>
</table>

element **TURN-TAKING**

<table>
<tr><td>Description</td><td colspan="5">Turn taking category.</td></tr>
<tr><td rowspan="3">Attributes</td><td>Name</td><td>Type</td><td>Use</td><td>Default</td><td>Description</td></tr>
<tr><td>TYPE</td><td><i>enum</i></td><td>required</td><td>-</td><td>The specific type.
Values [<i>take</i> | <i>give</i> | <i>keep</i> | <i>request</i> | <i>accept</i>].</td></tr>
<tr><td>ADDRESSEE</td><td><i>IDREF</i></td><td>required</td><td>-</td><td>Reference to the addressee of this function.</td></tr>
<tr><td rowspan="2">Elements</td><td>Name</td><td>Occurs</td><td colspan="3">Description</td></tr>
<tr><td>-</td><td>-</td><td colspan="3">-</td></tr>
</table>

element **SPEECH-ACT**

Description	Speech acts category.				
Attributes	**Name**	**Type**	**Use**	**Default**	**Description**
	TYPE	*enum*	required	-	The specific type. Values [*inform* | *ask* | *request*].
Elements	**Name**	**Occurs**	**Description**		
	-	-	-		

element **GROUNDING**

Description	Grounding category.				
Attributes	**Name**	**Type**	**Use**	**Default**	**Description**
	TYPE	*enum*	required	-	The specific type. Values [*request-ack* | *ack* | *repair* | *cancel*].
Elements	**Name**	**Occurs**	**Description**		
	-	-	-		

D.2.3 Performative Track Elements

Every FML chunk in this track can include **zero** or **more** `<performative-extension>` elements that currently encapsulate mixed content (text or performative functions). This is meant to be a place holder for future extensions of the performative track that will take care of the performative functions, perhaps with more detailed ad-hoc standard representation languages.

In the description of the elements contained in the `<performative-extension>` element, we suggest possible type of functions that these elements can describe.

element **PERFORMATIVE-EXTENSION**

Description	A performative extension element for the performative track chunks.				
Attributes	Name	Type	Use	Default	Description
	ADDRESSEE	*IDREF*	optional	-	Reference to the addressee of this function.
Elements	Name	Occurs	Description		
	DISCOURSE-STRUCTURE	0..*	*topic, segment, . . .*		
	RHETORICAL-STRUCTURE	0..*	*elaborate, summarize, clarify, contrast, emphasize, . . .*		
	INFORMATION-STRUCTURE	0..*	*rheme, theme, given, new, . . .*		
	PROPOSITION	0..*	*Any formal notation (e.g. "own(A,B)").*		

D.2.4 Mental State Track Elements

element **COGNITIVE-PROCESS**

Description	A participant's cognitive process.				
Attributes	Name	Type	Use	Default	Description
	WEIGHT-FACTOR	*decimal*	required	-	Weight of the process specified with the *type* attribute in the range [0..1].
	TYPE	*enum*	required	-	The type of cognitive process. Values [*think* | *remember* | *infer* | *decide* | *idle*].
Elements	Name	Occurs	Description		
	-	-	-		

element **EMOTION**

Description	A participant's emotion.				
Attributes	Name	Type	Use	Default	Description
	WEIGHT-FACTOR	*decimal*	required	-	Weight of the emotion specified with the *type* attribute in the range [0..1].
	TYPE	*enum*	required	-	The type of emotion. Values [*anger* | *disgust* | *embarrassment* | *fear* | *happiness* | *sadness* | *surprise* | *shame*].
	REGULATION	*enum*	optional	-	The regulation of the emotion. Values [*felt* | *fake* | *inhibit*].
	INTENSITY	*decimal*	optional	0.5	The intensity of the emotion in the range [0..1].
Elements	Name	Occurs	Description		
	-	-	-		

E

XSLT Transformation Rules

```
<?xml version="1.0"?>

<xsl:transform version="1.0"
  xmlns:xsl="http://www.w3.org/1999/XSL/Transform"
  xmlns:tns="http://cadia.ru.is/FMLSpecification"
  xmlns="http://www.bml-initiative.org/bml/bml-1.0"
  xmlns:fga="http://cadia.ru.is/FMLGreetingAgent" >

  <xsl:variable name="defaultDuration" select="0.1" />

  <!-- Apply an Identity Transform -->
    <xsl:template match="node()|@*">
        <xsl:copy>
            <xsl:apply-templates select="node()|@*"/>
        </xsl:copy>
    </xsl:template>

  <!-- Copy the matching BML elements as they appear adding the BML 1.0 namespace -->
  <xsl:template match="bml|constraint|synchronize|before|after|sync">
    <xsl:element name="{local-name()}" namespace="http://www.bml-initiative.org/bml/bml-1.0" >
            <xsl:copy-of select="attribute::*"/>
            <xsl:apply-templates select="node()|@*"/>
        </xsl:element>
    </xsl:template>

  <!-- Template for all FML functions not handled by these rules, transforms them into <wait> elements -->
  <xsl:template match="bml/speech-act|bml/grounding|bml/back-channel|bml/turn-taking|bml/performative-extension|bml/cognitive-process|bml/emotion">

    <xsl:call-template name="show-comment-template"/>
        <xsl:element name="wait" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id">
      <xsl:value-of select="@id"/>
      </xsl:attribute>

      <xsl:choose>
             <xsl:when test='string-length(@start)>0'>
             <xsl:attribute name="start"><xsl:value-of select="@start"/></xsl:attribute>
             </xsl:when>
      </xsl:choose>

      <xsl:attribute name="duration"><xsl:value-of select="$defaultDuration"/></xsl:attribute>
    </xsl:element>
    </xsl:template>

  <!-- Template for the "break-away" function of the "closing" category in the Interactional track -->
  <xsl:template match="bml/closing[@type = 'break-away']">

    <xsl:call-template name="show-comment-template"/>

    <xsl:element name="gaze" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze</xsl:attribute>
      <xsl:attribute name="start">
        <xsl:choose>
          <xsl:when test='string-length(@start)>0'> <xsl:value-of select="@start"/></xsl:when>
          <xsl:otherwise>0</xsl:otherwise>
        </xsl:choose>
      </xsl:attribute>
      <xsl:attribute name="end">1</xsl:attribute>
      <xsl:attribute name="influence">EYES</xsl:attribute>
      <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
      <xsl:attribute name="offsetAngle">45.0</xsl:attribute>
      <xsl:attribute name="offsetDirection">DOWNLEFT</xsl:attribute>
    </xsl:element>
  </xsl:template>
```

FML Listing E.1: XSLT Transformation rules used by the FML-to-BML Transformer (page 1 out of 8).

```xml
<!-- Template for the "farewell" function of the "closing" category in the Interactional track -->
<xsl:template match="bml/closing[@type = 'farewell']">

  <xsl:call-template name="show-comment-template"/>

  <xsl:variable name="floorIDVar" select="@tns:floorID" />
  <xsl:variable name="characterIDVar" select="../@characterId"/>

  <xsl:variable name="floorElementVar" select="/saiba-act/declarations/floors/floor[@tns:floorID=$floorIDVar]" />
  <xsl:variable name="participantElementVar" select="$floorElementVar/participant[@identikitRef=$characterIDVar]" />

  <xsl:variable name="participantAttitudeAffiliationVar" select="$participantElementVar/attitude/@affiliation" />
  <xsl:variable name="participantAttitudeTowardsVar" select="$participantElementVar/attitude/@towards" />

  <xsl:element name="gaze" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze</xsl:attribute>
    <xsl:attribute name="start">
    <xsl:choose>
      <xsl:when test='string-length(@start)>0'> <xsl:value-of select="@start"/></xsl:when>
      <xsl:otherwise>0</xsl:otherwise>
    </xsl:choose>
    </xsl:attribute>
    <xsl:attribute name="end"><xsl:value-of select="@id"/>_gaze:start + 3</xsl:attribute>
    <xsl:attribute name="influence">HEAD</xsl:attribute>
    <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
  </xsl:element>

  <xsl:if test="$participantAttitudeAffiliationVar = 'FRIENDLY' and $participantAttitudeTowardsVar = @addressee">

    <xsl:element name="faceLexeme" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id"><xsl:value-of select="@id"/>_face_1</xsl:attribute>
      <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze:start</xsl:attribute>
      <xsl:attribute name="end"><xsl:value-of select="@id"/>_face_1:start + 4</xsl:attribute>
      <xsl:attribute name="attackPeak"><xsl:value-of select="@id"/>:start + 2</xsl:attribute>
      <xsl:attribute name="overshoot">0</xsl:attribute>
      <xsl:attribute name="lexeme">LOWER_BROWS</xsl:attribute>
      <xsl:attribute name="amount">0.8</xsl:attribute>
    </xsl:element>

    <xsl:element name="faceLexeme" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id"><xsl:value-of select="@id"/>_face_2</xsl:attribute>
      <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze:start</xsl:attribute>
      <xsl:attribute name="end"><xsl:value-of select="@id"/>_face_2:start + 4</xsl:attribute>
      <xsl:attribute name="attackPeak"><xsl:value-of select="@id"/>:start + 2</xsl:attribute>
      <xsl:attribute name="overshoot">0</xsl:attribute>
      <xsl:attribute name="lexeme">LOWER_BROWS_MOUTH_CORNERS</xsl:attribute>
      <xsl:attribute name="amount">0.8</xsl:attribute>
    </xsl:element>

  </xsl:if>

  <xsl:element name="gesture" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_gesture</xsl:attribute>
    <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze:ready</xsl:attribute>
    <xsl:attribute name="lexeme">WAVE</xsl:attribute>
  </xsl:element>

</xsl:template>
```

FML Listing E.2: XSLT Transformation rules used by the FML-to-BML Transformer (page 2 out of 8).

```
<!-- Template for the functions of the "initiation" category in the Interactional track -->
<xsl:template match="bml/initiation">

  <xsl:call-template name="show-comment-template"/>

  <xsl:variable name="floorIDVar" select="@tns:floorID" />
  <xsl:variable name="characterIDVar" select="../@characterId"/>

  <xsl:variable name="floorElementVar" select="/saiba-act/declarations/floors/floor[@tns:floorID=$floorIDVar]"
       />
  <xsl:variable name="participantElementVar" select="$floorElementVar/participant[@identikitRef=$
       characterIDVar]" />

  <xsl:variable name="participantAttitudeAffiliationVar" select="$participantElementVar/attitude/@affiliation"
       />
  <xsl:variable name="participantAttitudeTowardsVar" select="$participantElementVar/attitude/@towards" />

  <xsl:variable name="participantPersonalityVar" select="/saiba-act/declarations/identikits/identikit[@id=$
       characterIDVar]/personality/@extraversion" />

  <!-- Option for the "react" function -->
  <xsl:choose>
    <xsl:when test="@type = 'react'">

      <xsl:element name="gaze" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/></xsl:attribute>
        <xsl:attribute name="start">
        <xsl:choose>
          <xsl:when test='string-length(@start)>0'>
          <xsl:value-of select="@start"/>
          </xsl:when>
          <xsl:otherwise>
          0
          </xsl:otherwise>
        </xsl:choose>
        </xsl:attribute>
        <xsl:attribute name="end"><xsl:value-of select="@id"/>:start + 1</xsl:attribute>
        <xsl:attribute name="influence">HEAD</xsl:attribute>
        <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
      </xsl:element>

    </xsl:when>

    <!-- Option for the "recognize" function -->
    <xsl:when test="@type = 'recognize'">
      <xsl:variable name="addresseeVar" select="@addressee" />
      <xsl:variable name="relationshipElementVar" select="/saiba-act/declarations/identikits/identikit[@id=$
           characterIDVar]/relationships/relationship[@with = $addresseeVar]"/>

      <xsl:element name="gaze" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze_1</xsl:attribute>
        <xsl:attribute name="start">
        <xsl:choose>
          <xsl:when test='string-length(@start)>0'>
          <xsl:value-of select="@start"/>
          </xsl:when>
          <xsl:otherwise>
          0
          </xsl:otherwise>
        </xsl:choose>
        </xsl:attribute>
        <xsl:attribute name="end"><xsl:value-of select="@id"/>_gaze_1:start + 3</xsl:attribute>
        <xsl:attribute name="influence">HEAD</xsl:attribute>
        <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
      </xsl:element>
```

FML Listing E.3: XSLT Transformation rules used by the FML-to-BML Transformer (page 3 out of 8).

```
    <xsl:if test="$relationshipElementVar/@level = 'ACQUAINTANCE'">

      <xsl:element name="head" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/>_headtoss</xsl:attribute>
        <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start + 0.6</xsl:attribute>
        <xsl:attribute name="end"><xsl:value-of select="@id"/>_headtoss:start + 0.8</xsl:attribute>
        <xsl:attribute name="lexeme">TOSS</xsl:attribute>
      </xsl:element>

      <xsl:element name="faceLexeme" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/>_raisebrows</xsl:attribute>
        <xsl:attribute name="start"><xsl:value-of select="@id"/>_headtoss:start</xsl:attribute>
        <xsl:attribute name="end"><xsl:value-of select="@id"/>_headtoss:end</xsl:attribute>
        <xsl:attribute name="attackPeak"><xsl:value-of select="@id"/>:start + 0.4</xsl:attribute>
        <xsl:attribute name="lexeme">RAISE_BROWS</xsl:attribute>
        <xsl:attribute name="amount">0.5</xsl:attribute>
      </xsl:element>

    </xsl:if>

    <xsl:if test="$relationshipElementVar/@level = 'FRIEND'">

      <xsl:element name="head" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id"><xsl:value-of select="@id"/>_headnod</xsl:attribute>
      <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start + 0.5</xsl:attribute>
      <xsl:attribute name="end"><xsl:value-of select="@id"/>_headnod:start + 0.8</xsl:attribute>
      <xsl:attribute name="lexeme">NOD</xsl:attribute>
      </xsl:element>

    </xsl:if>

  </xsl:when>
```

FML Listing E.4: XSLT Transformation rules used by the FML-to-BML Transformer (page 4 out of 8).

```
<!-- Option for the "distant salutation" function -->
<xsl:when test="@type = 'salute-distant'">
  <xsl:element name="postureShift" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_postureShift</xsl:attribute>
    <xsl:attribute name="start">
      <xsl:choose>
        <xsl:when test='string-length(@start)>0'>
        <xsl:value-of select="@start"/>
        </xsl:when>
        <xsl:otherwise>
        0
        </xsl:otherwise>
      </xsl:choose>
    </xsl:attribute>
    <xsl:element name="stance" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="type">STANDING</xsl:attribute>
    </xsl:element>
    <xsl:element name="pose" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="part">WHOLEBODY</xsl:attribute>
      <xsl:attribute name="lexeme">FACE</xsl:attribute>
      <xsl:attribute name="fga:target"><xsl:value-of select="@addressee"/></xsl:attribute>
    </xsl:element>
  </xsl:element>

  <xsl:element name="gazeShift" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze_1</xsl:attribute>
    <xsl:attribute name="start">
    <xsl:choose>
      <xsl:when test='string-length(@start)>0'>
      <xsl:value-of select="@start"/>
      </xsl:when>
      <xsl:otherwise>
      0
      </xsl:otherwise>
    </xsl:choose>
    </xsl:attribute>
    <xsl:attribute name="influence">HEAD</xsl:attribute>
    <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
  </xsl:element>
  <xsl:element name="gazeShift" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze_2</xsl:attribute>
    <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start + 2</xsl:attribute>
    <xsl:attribute name="influence">HEAD</xsl:attribute>
    <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
    <xsl:attribute name="offsetAngle">45.0</xsl:attribute>
    <xsl:attribute name="offsetDirection">DOWNLEFT</xsl:attribute>
  </xsl:element>

  <xsl:element name="gesture" namespace="http://www.bml-initiative.org/bml/bml-1.0">
    <xsl:attribute name="id"><xsl:value-of select="@id"/>_gesture</xsl:attribute>
    <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start + 1</xsl:attribute>
    <xsl:attribute name="lexeme">SHORT-WAVE</xsl:attribute>
    <xsl:attribute name="mode">RIGHT-HAND</xsl:attribute>
  </xsl:element>
```

FML Listing E.5: XSLT Transformation rules used by the FML-to-BML Transformer (page 5 out of 8).

```
      <xsl:if test="$participantAttitudeAffiliationVar = 'FRIENDLY' and $participantAttitudeTowardsVar = @addressee">


        <xsl:element name="faceLexeme" namespace="http://www.bml-initiative.org/bml/bml-1.0">
          <xsl:attribute name="id"><xsl:value-of select="@id"/>_face_1</xsl:attribute>
          <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start</xsl:attribute>
          <xsl:attribute name="end"><xsl:value-of select="@id"/>_face_1:start + 30</xsl:attribute>
          <xsl:attribute name="attackPeak"><xsl:value-of select="@id"/>:start + 8</xsl:attribute>
          <xsl:attribute name="overshoot">20</xsl:attribute>
          <xsl:attribute name="lexeme">RAISE_BROWS</xsl:attribute>
          <xsl:attribute name="amount">0.8</xsl:attribute>
        </xsl:element>

        <xsl:element name="faceLexeme" namespace="http://www.bml-initiative.org/bml/bml-1.0">
          <xsl:attribute name="id"><xsl:value-of select="@id"/>_face_2</xsl:attribute>
          <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze_1:start</xsl:attribute>
          <xsl:attribute name="end"><xsl:value-of select="@id"/>_face_2:start + 30</xsl:attribute>
          <xsl:attribute name="attackPeak"><xsl:value-of select="@id"/>:start + 8</xsl:attribute>
          <xsl:attribute name="overshoot">20</xsl:attribute>
          <xsl:attribute name="lexeme">RAISE_MOUTH_CORNERS</xsl:attribute>
          <xsl:attribute name="amount">0.8</xsl:attribute>
        </xsl:element>

      </xsl:if>
    </xsl:when>

    <!-- Option for the "approach react" function -->
    <xsl:when test="@type = 'approach-react'">
      <xsl:if test="$participantAttitudeAffiliationVar = 'FRIENDLY' and $participantAttitudeTowardsVar = @addressee">
        <xsl:element name="gaze" namespace="http://www.bml-initiative.org/bml/bml-1.0">
          <xsl:attribute name="id"><xsl:value-of select="@id"/></xsl:attribute>
          <xsl:attribute name="start">
          <xsl:choose>
            <xsl:when test='string-length(@start)>0'>
            <xsl:value-of select="@start"/>
            </xsl:when>
            <xsl:otherwise>
            0
            </xsl:otherwise>
          </xsl:choose>
          </xsl:attribute>
          <xsl:attribute name="end"><xsl:value-of select="@id"/>:start + 1</xsl:attribute>
          <xsl:attribute name="influence">HEAD</xsl:attribute>
          <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
        </xsl:element>
      </xsl:if>
    </xsl:when>

    <!-- Option for the "close salutation" function -->
    <xsl:when test="@type = 'salute-close'">
      <xsl:element name="gazeShift" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/>_gaze</xsl:attribute>
        <xsl:attribute name="start">
        <xsl:choose>
          <xsl:when test='string-length(@start)>0'>
          <xsl:value-of select="@start"/>
          </xsl:when>
          <xsl:otherwise>
          0
          </xsl:otherwise>
        </xsl:choose>
        </xsl:attribute>
        <xsl:attribute name="influence">HEAD</xsl:attribute>
        <xsl:attribute name="target"><xsl:value-of select="@addressee"/></xsl:attribute>
      </xsl:element>
```

FML Listing E.6: XSLT Transformation rules used by the FML-to-BML Transformer (page 6 out of 8).

```
    <xsl:if test="$participantPersonalityVar = 'HIGH'">
      <xsl:element name="postureShift" namespace="http://www.bml-initiative.org/bml/bml-1.0">
        <xsl:attribute name="id"><xsl:value-of select="@id"/>_postureShift</xsl:attribute>
        <xsl:attribute name="start"><xsl:value-of select="@id"/>_gaze:start</xsl:attribute>
        <xsl:element name="stance" namespace="http://www.bml-initiative.org/bml/bml-1.0">
          <xsl:attribute name="type">STANDING</xsl:attribute>
        </xsl:element>
        <xsl:element name="pose" namespace="http://www.bml-initiative.org/bml/bml-1.0">
          <xsl:attribute name="part">WHOLEBODY</xsl:attribute>
          <xsl:attribute name="lexeme">LEANING_FORWARD</xsl:attribute>
        </xsl:element>
      </xsl:element>
    </xsl:if>

  </xsl:when>

  <!-- Option for the "initiate" function -->
  <xsl:when test="@type = 'initiate'">

    <xsl:element name="gesture" namespace="http://www.bml-initiative.org/bml/bml-1.0">
      <xsl:attribute name="id"><xsl:value-of select="@id"/>_gesture</xsl:attribute>
      <xsl:attribute name="start">0</xsl:attribute>
      <xsl:attribute name="lexeme">OPEN-HANDS</xsl:attribute>
      <xsl:attribute name="mode">RIGHT-HAND</xsl:attribute>
    </xsl:element>

  </xsl:when>
</xsl:choose>

</xsl:template>
```

FML Listing E.7: XSLT Transformation rules used by the FML-to-BML Transformer (page 7 out of 8).

```
  <xsl:template match="turn-taking">
        <xsl:apply-templates/>
    </xsl:template>

  <!-- Template to show the original FML function element before the transformation in a comment in the BML block -->
  <xsl:template name="show-comment-template">
    <xsl:variable name="SpecialTagFullStringVar">
    <xsl:apply-templates select="." mode="serialize"/>
    </xsl:variable>

    <xsl:comment> tag [<xsl:value-of select="name(.)" />] floorID [<xsl:value-of select="@tns:floorID" />] ID [<xsl:value-of select="@id"/>] type [<xsl:value-of select="@type"/>] <xsl:if test="@addressee != ''">addressee [<xsl:value-of select="@addressee"/>] </xsl:if> track [<xsl:value-of select="@track"/>] start [<xsl:value-of select="@start"/>]  <xsl:if test="name(.) = 'performative-extension'"> performative-extension tag [<xsl:value-of select="$SpecialTagFullStringVar"/>] </xsl:if> <xsl:if test="@track = 'mental-state'"> <xsl:if test="name(.) = 'emotion'"> regulation [<xsl:value-of select="@regulation"/>] intensity [<xsl:value-of select="@intensity"/>]</xsl:if> weightFactor [<xsl:value-of select="@weightFactor"/>] </xsl:if> </xsl:comment>

  </xsl:template>

<!-- The following templates copy all element names and attributes not matched by previous template in the new document as they are -->
<xsl:template match="*" mode="serialize"><xsl:text>&lt;</xsl:text><xsl:value-of select="name(.)"/>&#xA0;<xsl:apply-templates select="@*" mode="serialize" /><xsl:text>></xsl:text><xsl:apply-templates mode="serialize"/><xsl:text>&lt;/</xsl:text><xsl:value-of select="name(.)"/><xsl:text>></xsl:text></xsl:template>

<xsl:template match="text()" mode="serialize"><xsl:value-of select="."/></xsl:template>

<xsl:template match="@*" mode="serialize">&#xA0;<xsl:value-of select="name()" />="<xsl:value-of select="." />"</xsl:template>

<!-- Template to show the declarations section as a comment at the beginning of the transformed document -->
<xsl:template match="declarations">
<xsl:comment>
Tag [declarations]
[___IDENTIKITS___]
<xsl:for-each select="identikits/identikit">
id [<xsl:value-of select="@id" />] name [<xsl:value-of select="@name"/>] gender [<xsl:value-of select="@gender"/>] <xsl:if test="personality">personality(extraversion) [<xsl:value-of select="personality/@extraversion" />]</xsl:if><xsl:for-each select="relationships/relationship"> relationship [<xsl:value-of select="@level" /> with <xsl:value-of select="@with" />]</xsl:for-each>
</xsl:for-each>
[___FLOORS___]
<xsl:for-each select="floors/floor">
floorID [<xsl:value-of select="@tns:floorID" />] floor-cfg [<xsl:value-of select="@floor-cfg"/>]
<xsl:for-each select="participant">
Tag [participant] identikitRef [<xsl:value-of select="@identikitRef" />] role [<xsl:value-of select="@role"/>] entity [<xsl:value-of select="@entity"/>] <xsl:if test="attitude">attitude(affiliation) [<xsl:value-of select="attitude/@affiliation" /> towards <xsl:value-of select="attitude/@towards" />]</xsl:if>
</xsl:for-each>
</xsl:for-each>
</xsl:comment>
</xsl:template>

</xsl:transform>
```

FML Listing E.8: XSLT Transformation rules used by the FML-to-BML Transformer (page 8 out of 8).

School of Computer Science
Reykjavík University
Menntavegi 1
101 Reykjavík, Iceland
Tel. +354 599 6200
Fax +354 599 6201
www.reykjavikuniversity.is
ISSN 1670-8539