

Detecting digital chameleons

Jeremy N. Bailenson^{a,*}, Nick Yee^a, Kayur Patel^b,
Andrew C. Beall^c

^a Department of Communication, Stanford University, Stanford, CA 94305, United States

^b Department of Computer Science, Stanford University, United States

^c Department of Psychology, University of California, Santa Barbara, United States

Available online 2 March 2007

Abstract

Conversations are characterized by an interactional synchrony between verbal and nonverbal behaviors [Kendon, A. (1970). Movement coordination in social interaction: some examples described. *Acta Psychologica*, 32(2), 101–125]. A subset of these contingent conversational behaviors is direct mimicry. During face to face interaction, people who mimic the verbal [Giles, H., Coupland, J., & Coupland, N. (1991). Accommodation theory: Communication, context, and consequence. In Giles, H., Coupland, J., & Coupland, N. *Contexts of accommodation. Developments in applied socio-linguistics*. Cambridge: Cambridge University Press] and nonverbal behaviors [Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76, 893–910] gain social advantage. Most research examining mimicry behavior in interaction examines ‘implicit mimicry’ in which the mimicked individual is unaware of the behavior of the mimicker. In this paper, we examined how effective people were at explicitly detecting mimicking computer agents and the consequences of mimic detection in terms of social influence and interactional synchrony. In Experiment 1, participant pairs engaged in a “one-degree of freedom” Turing Test. When the computer agent mimicked them, users were significantly worse than chance at identifying the other human. In Experiment 2, participants were more likely to detect mimicry in an agent that mirror-mimicked their head movements (three degrees of freedom) than agents that either congruently mimicked their behaviors or mimicked those movements on another rotational axis. We discuss implications for theories of interactivity.

© 2007 Elsevier Ltd. All rights reserved.

Keywords: Embodied agents; Social interaction; Virtual reality; Turing test

* Corresponding author. Tel.: +1 650 723 0701; fax: +1 650 723 2472.

E-mail address: Bailenson@Stanford.edu (J.N. Bailenson).

1. Introduction

Previous research has shown that people are uniquely influenced by others who mimic their language (Giles, Coupland, & Coupland, 1991; Van Baaren, Holland, Steenaert, & van Knippenberg, 2003) or their gestures (Chartrand & Bargh, 1999) during social interaction. Moreover, recent research has extended these findings to computer agents: voice synthesizers that mimic vocal patterns (Suzuki, Takeuchi, Ishii, & Okada, 2003) as well as embodied agents in immersive virtual reality that mimic nonverbal behavior (Bailenson & Yee, 2005). Given the growing prevalence of digital conversational agents in educational software (Graesser et al., 2003), internet applications (Yee, 2006), communications systems (Cassell, 2001), entertainment (Konijn & Hoorn, 2004), and devices such as cellular phones and personal data assistants, the potential use of algorithmic mimicry strategies in software agents is extraordinary. The goal of the current research is to (1) understand and explore how astute people are in detecting *digital chameleons*, agents who utilize mimicry algorithms, and (2) examine the implications of explicit mimicry detection in terms of social perceptions of the mimickers. While previous research has demonstrated the advantages of implicit mimicry, as far as we know this is the first study to report results from the explicit detection of various types of mimicry.

1.1. Interactional synchrony in discourse

In early research of verbal and nonverbal behavior, Kendon (1970) recorded and closely analyzed filmed interactions in slow motion. In these video analyses, Kendon noted and recorded every “minimally perceptible change” of each body part along a multi-layered time-line. These meticulous analyses revealed three kinds of synchrony. First, the points of change in the movement of separate body parts of the same individual coincided. Second, these changes in multiple body parts coincided with changes in speech of that individual. For example, larger body movements coincided with sentences and phrases, while smaller body movements coincided with syllabic or sub-syllabic changes. The final type of synchrony is *interactional synchrony*: The extent to which, in a dyadic interaction, the points of change of person A are in synchrony with the points of change in person B. Interactional synchrony was found to occur from the sub-syllabic level to the word level. In other words, synchrony at the phrase or sentence level was less consistent than at the word level or below.

Kendon argued that interactional synchrony functions as the regulator of the “delicate coordination of expectancies among participants,” (Kendon, 1970, p. 76) and suggested that interactional synchrony impacts credibility, persuasion, and trust in interactions by managing expectancies among participants. In essence, synchrony results when people’s behaviors are strategically contingent upon one another. One of the most obvious examples is direct mimicry – when one person directly repeats a verbal or nonverbal action in conversation.

1.2. Implicit mimicry in verbal and nonverbal behavior

It seems almost second nature that, without being aware of doing so, people mimic one another in social situations; this phenomenon occurs with laughter (Provine, 1992), eating habits (Johnston, 2002), mood (Neumann & Strack, 2000), and verbal behavior (Cappella

& Panalp, 1981; Giles, 1973; Levelt & Kelter, 1982). Dozens of studies have demonstrated communicators' implicit *convergence* – mutual adaptation of linguistic, prosodic, and non-verbal features (see Giles et al., 1991, for a review of this work).

Given the extremely high prevalence of mimicry in social interaction, researchers have sought to understand what causes this behavior. In terms of verbal behavior, accommodation is often performed in order to achieve strategic conversational goals, even if the mimickers remain unaware of their own accommodation. For example, Welkowitz, Feldstein, Finkelstein, and Aylesworth (1972) demonstrated that people in dyads who perceived themselves to be more similar converged more often in terms of vocal intensity than people who were randomly paired together. Furthermore, matching the speech rate of another person while requesting assistance was more likely to elicit positive responses than utilizing non-converging rates (Wiener & Mehrabian, 1968). Similarly, people who have a high need to gain approval from others are highly likely to converge (Natale, 1975).

In terms of nonverbal behavior, initial work by LaFrance (1979) demonstrated evidence of a correlation between nonverbal synchrony of posture mimicry and interaction quality. In classrooms where nonverbal gestures from the students were synchronized with gestures from the teacher, students reported a better rapport with the teacher than in classrooms in which the teacher's gestures were not mimicked by the students. Consequently, other researchers have explored the notion that people who utilize mimicry in interaction position themselves to gain advantage or to exert social influence in that interaction. Chartrand and Bargh (1999) discuss the *chameleon effect*, namely that people tend to unconsciously mimic the nonverbal behaviors of others in the context of social interaction. Furthermore, those researchers demonstrated that participants interacting with a confederate who mimicked their behavioral mannerisms liked that confederate more and rated the interaction quality higher than participants who interacted with confederates who displayed neutral, non-descript mannerisms. According to Chartrand and Bargh, "the chameleon effect operates in a passive, non-goal dependent manner to create greater liking and ease of interaction," (p. 901).

Subsequent work has replicated the unique social influence of verbal chameleons; for example, waiters who repeated customers' orders received larger tips than waiters who did not (Van Baaren et al., 2003). Moreover, a person who is mimicked tends to exhibit prosocial behaviors not only towards the mimicker but towards other people as well (Van Baaren, Holland, Kawakami, & van Knippenberg, 2004).

1.3. Explicit mimicry detection

In most of the research discussed above, nonverbal mimicry was kept implicit, in the sense that the people being mimicked had no conscious awareness of the mimicry. Furthermore, those who mimicked did so passively, in a non-goal oriented manner.

The difference between implicit and explicit knowledge received ample attention initially in terms of memory and brain structure (e.g., Graf & Schacter, 1985) as well as more recently in attitudes and social behaviors (see Dovidio, Kawakami, & Beach, 2001, for a review). In the current work, we examine people's ability to explicitly detect (i.e., have conscious awareness of) mimicry behavior. Since computer agents are necessarily programmed with defined goals, any mimicry that occurs is in some sense intentional. Consequently, an important question concerns the implications of mimicry detection for

the interaction: how will people react to an interactant when they know that interactant's behavior is a direct mimic of their own?

In terms of detecting mimicry in verbal behavior, research has shown that there are costs associated with high amounts of convergence when it is explicitly detected during communications. Bourhis, Giles, and Lambert (1975) demonstrated that convergence along multiple channels simultaneously such as information content, grammar, and vocabulary is associated with decreased trustworthiness according to third-person listeners. According to those authors, converging at an extremely high rate raises doubts about the social identity of the mimicker. According to Giles (1980), communicators have a threshold for tolerating convergence in conversation. Once that barrier is crossed, people create the impression that the mimicker is patronizing them and begin to feel uncomfortable. As such, the convergence behavior, designed to increase social integration, actually backfires when it is detected.

In terms of detecting nonverbal mimicry, to our knowledge, there is very little work on this issue. According to Chartrand (personal communication, February, 2005) one study on face-to-face interaction featured confederates who nonverbally mimicked participants who were explicitly instructed about the mimicry. The researchers conducting that study were forced to discontinue the research because participants being mimicked reported extremely high discomfort levels during the interaction. On the one hand, this high level of discomfort makes sense – children taunt one another by parroting speech and intonation. However, why would implicit awareness of mimicry produce social influence and explicit awareness produce extreme discomfort? While there is often a very low correlation between implicit and explicit attitudes (Dovidio et al., 2001), such opposing reactions are puzzling.

If one conceives intentional mimicry as a form of deception, then explicit detection of mimicry should evoke a negative response. Various theoretical approaches outline the motivations and consequences of deception in human communication, such as interpersonal deception theory (Buller & Burgoon, 1996) or Zuckerman's model of nonverbal leakage (Zuckerman, DePaulo, & Rosenthal, 1981). Common to these theories are two assumptions: that people sometimes deceive in order to gain advantage, and that advantage is lost when the deception is detected. It is doubtful that nonverbal mimicry is always a form of intentional deception in face-to-face interaction. However, in this case, what is important is the subjective interpretation of the person being mimicked – if he or she perceives the behavior to be an intentional ploy by a person or by a computer program to gain advantage, then the response will likely be negative. In the current work, we can examine this hypothesis by comparing the response towards agents from people who detected the mimic to the response from those who did not. Given that the discourse we study in the current work is in the context of attitude change, the deception framework is quite relevant.

1.4. Automated mimicry in computer-mediated communication

Nonverbal chameleons need not be human beings. When humans process narrative text containing other characters, the distinction between self and other becomes blurry, with the human assimilating and absorbing various traits and features of the fictional characters (Kuiken et al., 2004). Interaction with digital representations of people further demonstrates this assimilation; recent work has suggested that social influence resulting

from mimicry is especially powerful with embodied computer agents. [Bailenson and Yee \(2005\)](#) examined situations in which participants interacted with an embodied agent inside of an immersive virtual reality simulation. In that research, the realistic-looking agent administered a three-minute verbal persuasive message. For half of the subjects, the agent was a digital chameleon, in that its head movements were an exact mimic of the subject's head movements at a four second delay. For the other half, the agent's head movements were a mimic of a different subject's head movements. Results demonstrated two important findings: (1) participants rarely (i.e., 5% of all instances) detected their own gestures when those gestures were utilized by the agents at a four second delay, and (2) participants were more persuaded by the agent, liked the agent more, and actually looked at the agent more in the mimic condition than in the recorded condition. In addition, research on verbal behavior has demonstrated that a computer agent whose voice mimics properties of a human being's voice is particularly effective compared to other types of agent vocal behavior ([Suzuki et al., 2003](#)).

In computer-mediated communication, people interact with one another through digital representations, whether that representation is a voice on a digital cell phone, a stream of text in a chat room, an image on a videoconference, or an avatar in a video game or virtual reality simulation. For example, [Shamp \(1991\)](#) found evidence that people who exchanged lots of personal information prevented themselves from being "mechanomorphized", that is, perceived by others computer-like. Similarly, using digital chameleons in this media remains a distinct possibility, one that is in many ways more powerful than face-to-face mimicry. For example, in computer-mediated communication, algorithms can be indiscriminately and constantly applied to mimic peoples' behaviors. Once mimics are utilized on an algorithmic level, the potential to utilize more fine grained strategies such as probabilistic mimics and scaled mimics become almost trivial to employ. Such a systematic application of various levels of mimicry is difficult to implement in face-to-face conversation compared to the ease with which it can be employed in a digitally mediated setting.

Previous research discusses the notion of Transformed Social Interaction (TSI; [Bailenson, Beall, Loomis, Blascovich, & Turk, 2004](#); [Bailenson & Beall, 2006](#)). TSI allows users of computer-mediated communication to strategically decouple their rendered behavior (i.e., the behaviors that other people see) from their actual behavior (i.e., the behaviors that were actually performed). One of the most intuitive applications of TSI is nonverbal mimicry. Computer-mediated communication systems, most notably immersive virtual reality simulations that track and render behaviors from multiple users simultaneously, *necessarily* archive extremely detailed records of all verbal and nonverbal behaviors. In other words, in order to make the interaction function properly (i.e., show the behaviors of one interactant to another), the system needs to collect all actions from all participants ([Blascovich et al., 2002](#)). Consequently, it is possible for an interactant seeking social influence to use simple algorithms to automatically mimic the behaviors of any number of other interactants.

1.5. Digital mimicry and discourse content

It is clear that digital mimicry will be more effective and more easily detected in certain discourse structures and contents. A literature review that compared face-to-face interaction to computer-mediated communication ([Whittaker, 2002](#)) indicated that visual,

nonverbal information in computer-mediated communication is uniquely critical for tasks where affect and attitude play a central role (Kiesler, Siegel, & McGuire, 1984; Morley & Stephenson, 1969; Short, Williams, & Christie, 1976). In other words, in situations in which interactants discuss social information, nonverbal behavior is a necessary and diagnostic cue. However, when the discourse topic is largely task oriented, nonverbal behavior is extraneous. This distinction is quite relevant to the current work, which utilized discourse centered around attitude change. Given the importance of visual information in these social and persuasive tasks, interactants should be highly susceptible to mimicry and accommodation via head movements.

On the other hand, in certain interactions using mimicry as a strategy may backfire. When the verbal content in discourse clearly cues that the mimicker is attempting to gain social advantage, such as portraying a particular economic goal (Cohen & Cooper, 1986), it is likely that the use of automated mimicry will be detected more often and will consequently be less effective. In the following section, we discuss the pervasiveness of computer-mediated communication, and the many ways in which digital chameleons can exist in current media.

1.6. Current digital human representations

Recent estimations show that 53 million people in the United States use instant messaging (IM; [Pew Internet & American Life Project, 2004](#)). Among those born in 1977 or after, 62% use instant messaging on a regular basis. Not only is IM used for communication among friends and family (33% of all IM users), but the same report shows that the majority of IM users use the tool for collaboration with co-workers (40% of all IM users). The digital format of this communication channel allows mimicry on many levels. For example, the IM client can match the typing speed and response delay of interactants. The client can also offer summaries of the average sentence length, use of punctuation and other syntactic structures of the conversation. The same mimicry can also be applied to other collaborative environments (e.g., [Avouris, Dimitracopoulou, & Komis, 2003](#)) as well as more asynchronous environments such as message boards and email where algorithms can be used to calculate how closely a reply matches a desired email or forum message in style and structure.

More elaborate mimicry can be achieved in environments where visually embodied avatars or agents exist, such as in online social environments where thousands of users interact concurrently through visually embodied avatars. In one type of these environments, known as massively-multiplayer online role-playing games (MMORPGs), millions of users spend on average 22 h a week interacting with each other through avatars ([Yee, 2006](#)). In these environments, the system can have agents mimic an avatar's movement patterns, clothing, facial features, height, weight, hair style, or perhaps even adopt a name that starts with the same first letter as the avatar's. The system can also track a user's preferred greeting phrases and have agents greet the user with a probabilistic sampling of those phrases. Moreover, given the huge advances in computer vision and algorithmic gesture detection ([Hu, Chang, Feris, & Turk, 2004](#)), coupled with the propensity for more and more computers to utilize high-bandwidth connections and embedded video cameras, the potential for computer agents to detect, mimic, and implement human gestures (e.g., [Fussell et al., 2004](#)) and other behaviors is quite boundless.

1.7. Overview of experiments

Two experiments examined people interacting with mimicking algorithms. The first experiment examined people's ability to differentiate between real people and digital chameleons. The second study examined different types of digital chameleons that varied in levels of interactivity and tested the effect of mimicry detection on social influence.

In Experiment 1, participants engaged in a variation of the Turing Test (Turing, 1950) on a desktop computer. Pairs of interactants communicated with one another by pushing a button and seeing two indicators, one of which lit up when the other participant hit his or her button, and the other of which lit up according to a computer algorithm. Participants attempted to determine which of the two buttons represented the actual human and rated their confidence in that decision.

In Experiment 2, participants entered an immersive virtual environment and listened to a persuasive message administered by an embodied agent. We varied the veridicality of the mimic behavior. The agent either *mirror-mimicked* the participants' head movements (left was left), *congruently-mimicked* their head movements (left was right), or *axis-shift-mimicked* their head movements transformed onto a separate axis as it verbally delivered a persuasive message (left was up). Participants rated the embodied agent on a number of subjective measures and wrote open-ended responses about the agent. We measured how often they detected the mimic as well as how effective they rated the embodied agent as an interactant.

2. Experiment 1

In Experiment 1, two participants engaged in a “one degree of freedom” Turing Test by communicating with another human being with the space bar and viewing two indicators: one that lit up instantaneously when the other participant hit the space bar, and another that responded according to a computer algorithm. We manipulated the type of algorithms utilized by the computer. Given previous findings discussed above that indicate people have a difficult time noticing mimicry during interaction, we predicted that people would fail the Turing Test more often with mimicking algorithms than other types of algorithms.

3. Method

3.1. Design

We manipulated two within-subjects variables: *agent behavior* (mimic, recorded, random, pattern, and mimic/pattern) and *noise factor* (the amount of random variance that was added or subtracted to the duration of each button press and release – 0% or .25%). In the *mimic* condition, the agent mimicked the button pressing of the user at a four second lag. In the *recorded* condition, we played back one of nine recorded data files from three human beings who performed the task during pilot testing. In the *random* condition, the computer generated random noise and toggled the button back and forth according to that noise. In the *pattern* condition, the computer randomly picked between two simple preprogrammed patterns (either Morse code for SOS or a sequence that toggled between on and off, both steadily increasing and decreasing according to an exponential algorithm). In the *mimic/pattern* condition, for 50% of the time the computer mimicked the




Fig. 1. A screenshot of the participants' task in Experiment 1.

human at a four second lag, and for the other 50% picked one of the two patterns. These particular algorithms were chosen as simple prototypes for various types of classes of human behavior emulation.

For each participant there were 16 trials: four trials of mimic (two with each level of noise factor), four trials of recorded (two with each level of noise factor), four trials of pattern (two with each level of noise factor), two random trials, and two trials of mimic/pattern (one with each level of noise factor).¹ Each participant viewed these 16 experimental trials in a random order after receiving two practice trials in which they familiarized themselves with the interface.

3.2. Materials

The interface consisted of a desktop computer screen that was separated into three distinct panels, as Fig. 1 demonstrates. Two of the panels, the left and right, were communication panels. These panels were randomly assigned as either a human or a computer player. Players communicated with one another by pressing the space bar. When they did so, the other player saw one of the indicator circles change color (green when the button was hit and red when it was not). Players and the computer communicated both by controlling the state of the circle (on or off) and the duration (how long the circle was on or off). Furthermore, the panel in the middle lit up when participants hit their own space bar, in order to provide them further information about their actions.

3.3. Procedure

Participants came into the lab in pairs, and sat in different rooms from one another. They were instructed to accomplish two goals: (1) use the space bar as effectively as possible to communicate their humanity to the other participant, and (2) to determine in each

¹ We did not completely cross the noise factor with the agent behavior factor because it did not make sense to add noise to a random pattern. Furthermore, the number of trials in the mimic/pattern condition was different from the others due to time limitations during the study, but tests of homogeneity of variance indicated that the variances in all conditions were not different from one another.

trial which of the two indicator circles (the left or the right) was the other human being. The other human was randomly placed on either the right side or the left side for each trial. Each trial lasted 60 s, after which the screen shown in Fig. 1 disappeared and a response screen appeared. Participants then chose which indicator they believed was controlled by the other human (left or right) and indicated how confident they were in their response on a scale that ranged from 1 (not confident at all) to 7 (extremely confident).

3.4. Participants

Participants were 26 undergraduate students (four female pairs, four male pairs, and five mixed pairs) who were paid for their participation.

4. Results and discussion

To analyze the data, we computed two dependent variables, percentage correct (i.e., correctly choosing the other human) and confidence. We first examined the differences between agent behaviors. Fig. 2 demonstrates the mean percentage correct in the five conditions.

We ran a within-group ANOVA with group (i.e., pair of subjects) as the random factor,² agent behavior as the independent variable, and percent correct as the dependent variable. There was a significant effect of agent behavior, $F(4, 48) = 3.85$, $p < .01$, partial $\eta^2 = .24$. Furthermore, post-hoc Tukey's HSD tests indicated that, with an alpha level of .05, participants performed significantly worse in the mimic condition than in all other conditions. No other differences were significant. Participants in the mimic condition performed significantly worse than chance ($t(12) = -2.21$, $p < .05$), while participants in the pattern ($t(12) = 2.89$, $p < .01$) and in the recorded condition ($t(12) = 2.55$, $p < .05$) performed significantly better than chance. Participants in the mimic/pattern condition ($t(12) = 2.55$, $p < .05$) and the random condition ($t(12) = 2.55$, $p < .05$) were not significantly different from chance.

We next ran an analysis to examine the simultaneous effects of agent behavior and noise factor. In order to do so, we did not include trials from the random condition in this analysis, because it was not possible to cross that condition with noise factor. We ran a 2×4 within-groups ANOVA with agent behavior (four levels: mimic, recorded, pattern, mimic/pattern) and noise factor (two levels: no noise or 25% noise) as independent variables and percentage correct as the dependent variables. Similar to the previous analysis, there was a significant effect of agent behavior, $F(3, 36) = 4.96$, $p < .01$, partial $\eta^2 = .29$. However, neither the effect of noise factor, $F(1, 12) = 1.32$, $p < .27$, partial $\eta^2 = .10$, nor the interaction, $F(3, 36) = .38$, $p < .86$, partial $\eta^2 = .03$ was significant. Furthermore, post-hoc Tukey's HSD tests indicated that, with an alpha level of .05, participants performed significantly worse in the mimic condition than in all other conditions. No other differences were significant.

We next examined confidence ratings, running a within-group ANOVA with group as the random factor, agent behavior as the independent variable, and mean confidence

² In other words, each independent observation was the average of all trials from both subjects in a given condition. We chose to do the analysis conservatively by-group instead of by-subject in order to preserve independence among the observations in the ANOVA.




Fig. 2. Average percent participants correctly identified the human by condition.

rating as the dependent variable. There was a significant effect of agent behavior, $F(4, 48) = 2.71, p < .05$, partial $\eta^2 = .18$. To further investigate the simultaneous effect of agent behavior and noise on confidence ratings, we ran a 2×4 within-groups ANOVA with agent behavior (four levels: mimic, recorded, pattern, pattern/recorded) and noise factor (two levels: no noise or 25% noise) as independent variables and confidence ratings as the dependent variables. As Fig. 3 demonstrates, there was a significant effect of agent behavior, $F(3, 36) = 3.57, p < .05$, partial $\eta^2 = .23$. However, neither the effect of noise factor, $F(1, 12) = .61, p < .45$, partial $\eta^2 = .05$, nor the interaction, $F(3, 36) = .88, p < .46$, partial $\eta^2 = .07$ was significant. Post-hoc Tukey’s HSD tests indicated that, with an alpha level of .05, participants were more confident in the mimic condition than in all other conditions. No other pairwise differences were significant. While the interaction between the two factors was not significant, it is interesting to note that the recorded condition is the only one in which adding noise results in lower confidence than the other conditions. It could be the case that direct recordings are unique in terms of the ability to implement simple alterations using computer algorithms.




Fig. 3. Average confidence ratings by noise and agent behavior variables in Experiment 1.

In sum, when an agent mimics a human participant, not only was that agent more effective than other types of agents at masquerading as a human, but participants actually performed worse than chance when attempting to discriminate the mimicker from the human being. Furthermore, they were most confident in their decision during mimic trials, despite the fact that they were actually performing the worst in that condition. This finding resonates with other research demonstrating that people often tend to be quite overconfident in wrong decisions (e.g., Kahneman & Tversky, 1973; Loftus, 1996).

In this study, people regarded implicit mimicry as a prototypical human trait during interaction. Mimickers were more likely to pass the Turing Test than actual humans! Nothing in the natural world mimics perfectly – confederates in experiments and interactants subconsciously conveying positive regard implement mimic behaviors that are only an approximation of the original behavior. However, digital chameleons mimic perfectly in terms of scope, scale and timing of any behavior. Consequently, if one takes as a given that people regard mimicking as a more favorable human trait than simple patterns or recordings, then it follows that a computer that mimics better than a human can is judged to be more human than the actual human. Of course this explanation is ad hoc, but previous work on synchrony and mimicry certainly provide some support for these findings.

In Experiment 2, we explored this mimicry effect further by examining agents' head movements in immersive virtual reality simulations. Head movements involve more degrees of freedom than simple button pushes (three compared to one). Consequently, we were able to examine agents that were contingent and interactive without exhibiting perfect mimicry. Furthermore, given that three degrees of freedom allows for more complex behavior, it should be more difficult for a mimicker to successfully imitate a human. Consequently, we can examine predictions about the effect of mimicry detection on subjective perception of the agent and social influence.

5. Experiment 2

Participants entered an immersive virtual environment and listened as an embodied agent (i.e., a human-like digital representation controlled by a computer) read them a persuasive message. The agent either exactly mirrored participants' head movements (mirror-mimic), reverse-mirrored the movements along the same axis (congruent-mimic), or mirrored the movements along a different axis (axis-switch). We predicted that the mimicry would be easiest to detect in the mirror-mimic condition because people see mirror reflections of themselves on a daily basis whereas the other two conditions do not occur in nature. Furthermore, as discussed in Section 1, we predicted that participants who did not detect the mimic would evaluate the agent more favorably in terms of trait ratings and agreement with a persuasive message than participants who explicitly detected the behavior.

6. Method

6.1. Design

We manipulated one between-subjects variable: *agent behavior* (mirror-mimic, congruent-mimic, or axis-shift-mimic). In the mirror-mimic condition, the agent's head movements exactly mimicked those of the participant at a four second delay on three axes

(pitch, yaw, and roll). In the congruent-mimic condition, the agent's head movements were exactly opposite of the participant on the same axis at a four second delay, such that if the participant looked ten degrees up or to the left, the agent would look ten degrees down or to the right. In the axis-switch condition, movements were directly switched from pitch to yaw and from yaw to pitch, such that if the participant looked up, the agent looked to the right, and if the participant looked to the left, the agent looked down. In all three conditions, we did not render movements greater than 90° on the agent, primarily to prevent the agent's chin from entering its chest when the participant moved his or her head very far to the left in the axis-switch condition.

Fig. 4 shows the male and female three-dimensional models utilized as agents. The agent was always the same gender as the participant. In both behavior conditions, the agent blinked (randomly according to a computer algorithm based on human blinking behavior) and exhibited lip movements driven by the amplitude of the recording of the persuasive passage. Other than lip movements, head movements, speech, and eye-blinks, there were no other behaviors exhibited by the agent.

6.2. Materials

In the virtual environment, participants were seated at a table, facing an embodied agent. Participants were able to see the head and shoulders of the agent. The same androgynous body was used for male and female agents.

6.3. Equipment

Fig. 5 shows a participant wearing the equipment needed to run the experiment. Participants wore a Virtual Research V8 stereoscopic head mounted display (HMD) that featured dual 680 horizontal by 480 vertical pixel resolution panels that refreshed at 60 Hz. or an nVisor SX HMD that featured dual 1280 horizontal by 1024 vertical pixel resolution panels that refreshed at 60 Hz. On both types of HMD, the display optics presented a visual field subtending approximately 50° horizontally by 38° vertically.



Fig. 4. The digital models of the male and female faces used in Experiment 2.

Perspectively correct stereoscopic images were rendered by a 1700 MHz Pentium IV computer with an NVIDIA GeForce FX 5200 graphics card, and were updated at an average frame rate of 60 Hz. The simulated viewpoint was continually updated as a function of the participants' head movements. The orientation of the participant's head was tracked by a three-axis orientation sensing system (Intersense IS250, update rate of 150 Hz). The system latency, or delay between a participant's head movement and the resulting concomitant update in the HMD's visual display was 45 ms maximum. The software used to assimilate the rendering and tracking was Vizard 2.15. Participants used a Logitech RumblePad Pro game pad to interact with the virtual environment.

6.4. Procedure

When participants arrived individually at the laboratory, an experimenter told them the following:

In this experiment, we are testing to see how people interact with other people inside virtual reality. The other virtual person in the room controlled by the computer is the presenter. The virtual person will present some information to you about a potential change in the security policy here at Stanford. We then will ask questions about the material he or she presented as well as some questions about the presenter.

Next, the experimenter instructed the participants on how to wear and adjust the HMD and how to use the game pad depicted in Fig. 5 to answer Likert-Scale questions. Once immersed, participants found themselves seated at a table directly across from an embodied agent. Depending on the assigned condition, the embodied agent varied in gender and behavior. The gender of the agent was set to match the gender of the participant. Once the participant was situated, the virtual agent delivered a persuasive message (based on work




Fig. 5. A depiction of our immersive virtual environment system. The components are: (1) audio output device, (2) HMD, (3) game pad input device, (4) image generator, and (5) orientation tracking device.

by Guadagno, Bailenson, Beall, Dimov, & Blascovich, 2005) about a campus security policy to make it mandatory for students to always carry their identification. The female agent delivered the message in a recorded female voice, and the male agent delivered the message in a recorded male voice. The length of each voice recording was 195 s.

After the agent finished delivering the recorded message, the agent disappeared and a blank screen filled the participant's view. The participant then received four statements (see the first four questions in the Appendix), one at a time, with which they indicated their agreement using a fully-labeled 7-point Likert-Scale from strongly disagree to strongly agree, numbered from negative three to positive three. They used the game pad depicted in Fig. 5 to register their responses. Next, participants removed the HMD and used pen and paper to answer the agent impression scales depicted in the Appendix. There were two reasons for asking the first four questions inside the virtual environment and the rest on pen and paper: (1) one of the three conditions of the current study was a replication of previous work and we wanted to keep the methodologies the same (Bailenson & Yee, 2005), and (2) keeping participants inside the virtual environment for too long risks simulator sickness, and given that the trait questionnaires were longer than the agreement questions we decided to present them on pen and paper. Finally, the participants wrote four open-ended paragraphs concerning their experience, the interaction, and the presenter's head movements. These questions appear in the Appendix.

6.5. Participants

Participants were recruited from an introductory communication class for course credit. There were 21 females (7 in the mirror-mimic condition, 7 in the congruent-mimic condition, and 7 in the axis-switch condition) and 25 males (8 in the mirror-mimic condition, 8 in the congruent-mimic condition, and 9 in the axis-switch condition). Experimental condition was randomly assigned.

6.6. Measures

6.6.1. Trait and agreement ratings

We ran a factor analysis to segment the questions on trait ratings of the agent (developed by Guadagno & Cialdini, 2002) as well as agreement questions about the content of the interaction depicted in the Appendix. These measures are discussed in Section 7.

6.6.2. Mimic detection

Three independent raters, blind to experimental condition, read the four open-ended paragraph responses from each participant, and each rater gave a one to participants they believed detected the mimic algorithm, and a zero for participants they believed did not detect the mimic algorithm. The inter-rater reliability for this measure was perfect.³

³ We chose to ask participants to write a paragraph (as opposed to a checklist of behaviors) because pilot studies have shown a huge tendency for participants to exhibit demand characteristics. Even when participants clearly had no idea the agent was mimicking them, very rarely would they admit to not noticing this behavior when presented in a recognition scale. While previous work (Van Baaren et al., 2003) has demonstrated the validity of checklists, it may be the case that people are embarrassed not to have noticed that a simple computer program has mimicked them.

7. Results and discussion

A principal components analysis on the 12 trait items and four agreement items revealed four factors that together accounted for 69% of the total variance. All factors had eigenvalues of over 1. The factor loadings of each item and the reliability of each factor are shown in Table 1. The first factor can be described as trustworthiness because it included honest, credible, competent, and sincere. The second factor can be described as warmth because it included friendly, warm, likeable, and approachable. The third factor can be described as informative because it included interesting, informed, and modest. The fourth factor can be described as agreement because it included the four questions about participants' agreement with the passage content. The scores for each factor were calculated for each participant using a regression method and thus were standardized.

We first ran an ANOVA with two between-subject independent variables (agent behavior and *detect*: whether or not the given participant detected the mimic algorithm or not) and presenter trustworthiness as the dependent variable. Gender was not included in this analysis because of low cell size given the factorial design resulting from three independent variables. There was a significant main effect of detect, $F(1,41) = 6.02, p = .04, \eta^2 = .09$. As Table 2 demonstrates, participants rated the presenter as less trustworthy when they detected the mimic ($M_{\text{yes}} = -.35, M_{\text{no}} = .39$). The effect of agent behavior was not significant ($F(2,41) = .33, p = .72, \eta^2 = .003$). The interaction effect was also not significant ($F(2,41) = .47, p = .63, \eta^2 = .02$).

We then ran a similar ANOVA with presenter warmth as the dependent variable. There was a significant main effect of detect, $F(1,41) = 7.20, p = .03, \eta^2 = .11$. Participants rated the presenter as less friendly when they detected the mimic ($M_{\text{yes}} = -.38, M_{\text{no}} = .43$). The effect of agent behavior was not significant ($F(2,41) = .08, p = .93, \eta^2 = .003$). The interaction effect was also not significant ($F(2,41) = .70, p = .50, \eta^2 = .03$).

Table 1
Factor loadings and reliability of factor items

Factor	Item	Factor loading
Trustworthiness $\alpha = .79$	Honest	.83
	Credible	.75
	Competent	.74
	Sincere	.71
	Trustworthy	.64
Warmth $\alpha = .87$	Friendly	.89
	Warm	.87
	Likeable	.82
	Approachable	.79
Informative $\alpha = .70$	Interesting	.81
	Informed	.78
	Modest	.73
Agreement $\alpha = .75$	Agreement	.91
	Valuable	.82
	Workable	.66
	Needed	.60

Table 2

Means and standard deviations of mimic conditions by mimic detection for presenter trustworthiness, warmth, informativeness, and agreement

	Mimic		Congruent		Axis-switch	
	No detect (<i>n</i> = 3)	Detect (<i>n</i> = 12)	No detect (<i>n</i> = 10)	Detect (<i>n</i> = 6)	No detect (<i>n</i> = 9)	Detect (<i>n</i> = 7)
Trustworthiness	.30 (.61)	−.15 (.86)	.41 (1.18)	−.88 (1.97)	.40 (1.19)	−.23 (.92)
Warmth	.48 (.70)	−.22 (1.00)	.19 (1.39)	−.28 (1.00)	.68 (1.37)	−.74 (1.04)
Informative	.29 (.67)	−.17 (1.05)	−.02 (1.29)	−.29 (1.32)	.14 (1.50)	.27 (1.25)
Agreement	−.70 (2.11)	.13 (1.32)	−.03 (1.52)	−.24 (.75)	.43 (.57)	−.22 (1.14)

We also ran an ANOVA with presenter informativeness as the dependent variable, but there were no significant main effects of agent behavior ($F(2,41) = .33, p = .72, \eta^2 = .02$) or detection ($F(1,41) = .25, p = .62, \eta^2 = .006$). The interaction effect was also not significant ($F(2,41) = .19, p = .83, \eta^2 = .009$).

Finally, we ran an ANOVA with agreement as the dependent variable, but there were no significant main effects of agent behavior ($F(2,41) = .32, p = .73, \eta^2 = .01$) or detection ($F(1,41) = .00, p = .99, \eta^2 = .00$). The interaction effect was also not significant ($F(2,41) = 1.11, p = .34, \eta^2 = .05$).

Next, we ran an ANOVA with gender and agent behavior as independent variables and detection (1 for yes, 0 for no) as the dependent variable. There was a significant effect of agent behavior, $F(2,41) = 3.22, p < .05, \eta^2 = .06$. Post-hoc Tukey's HSD tests ($\alpha = .05$) indicated that participants were more likely to detect the mirror-mimic condition (80%) than either the congruent-mimic (44%) condition or the axis-switch condition (38%). The pairwise difference between congruent-mimic and axis-switch was not significant, and neither was the main effect of gender ($F(1,41) = .17, p < .69, \eta^2 = .002$) nor the interaction between gender and agent behavior ($F(2,41) = .12, p < .89, \eta^2 = .002$).

The amount of mimic detection in the mirror-mimic condition, in the current study (80%) was quite a bit higher than a similar condition in our previous work (5%, [Bailenson & Yee, 2005](#)). The current study and the previous one were run with the same subject pool less than three months apart. Due to debriefing from the previous study, participants became aware of the manipulations through word of mouth, as some of the participants from the current study indicated during post-experiment informal interviews. Clearly this is a problematic finding, given that the purpose of the current work was to examine the detection of mimicry. While it is clearly the case that meaningful differences between the types of mimic behaviors exists (i.e., even with the word of mouth, participants still had trouble detecting the indirect mimics), the high baseline is alarming. Given that this one cell of our factorial design was a replication of previous findings, we wanted to ensure that the previous findings were still valid.

In order to partially remedy this shortcoming, six months after the completion of the current study (i.e., a new academic quarter), we ran a post-test that exactly replicated the mirror-mimic condition in a post-test with 53 paid participants (30 females) from a different academic year than the previous work. Only six of the naïve participants (11%) detected the mimic, a finding that closely resembles our initial work in this area.

In sum, in Experiment 2 an agent that mimics a user was seen as more effective if the user did not explicitly detect the mimic. Explicit detection of the mimic caused the agent to be evaluated poorly in terms of trustworthiness and the warmth factors, which is similar

to work by Bourhis et al. (1975) which demonstrated decreases in perceived trustworthiness for explicit convergence in verbal communication. However, the agent did not suffer in terms of more objective and substantive attributions such as informativeness and agreement. We had originally predicted a negative effect on all measures relating to both the agent and the message delivered by the agent. Instead, the results demonstrated that, while participants clearly punished the messenger for strategic nonverbal manipulation, they were able to dissociate the content of the message from the messenger.

In addition, mirror-mimics were detected more than the other types of mimics, which were detected less than half of the time, even with the contaminated subject pool. As the mimic behavior become less direct and more abstract, participants had a more difficult time detecting the behaviors, even when they were perfectly contingent.

8. General discussion

These two studies have examined interactants in discourse with various types of mimicking computer agents. In Experiment 1, many participants failed a “one degree of freedom” Turing Test with mimicking agents; they were actually more likely to believe the agent was a human than they would believe an actual human being was a human. Furthermore, they were extremely confident in their incorrect decisions when interacting with mimicking agents. In Experiment 2, with regards to head movements, participants were more likely to detect a mirror-mimic than either a congruent-mimic or an axis-switch-mimic. Furthermore, agents whose mimicry was undetected were more effective in terms of social personality ratings than agents whose mimicry was detected.

8.1. Rethinking the Turing test – verbal or nonverbal?

In the early 1950s, Alan Turing envisioned a test for artificial intelligence called the *Imitation Game*. The scenario consists of three participants: a human interrogator, a human player, and a computer player. Players attempt to convince the interrogator of their gender and the interrogator must choose one player as the human. Turing felt that creating an accurate physical simulation was too difficult and ultimately irrelevant for discerning intelligence. Therefore, he limited communication in the Imitation Game to a chat-like interface. The question for determining intelligence became, “Can Machines communicate in natural language in a manner indistinguishable from that of a human being?”

Every year researchers compete for the Loebner Prize, which offers a large sum of money to a designer who can design a natural language processing and generation system to fool a judge into believing a computer is human. As French (2000, p. 9) claims when reviewing the progress made in the first 50 years of attempts to pass the Turing Test, “To date, nothing has come remotely close to passing an unrestricted (Loebner) test and, as Dennett, who agreed to chair the event for its first few years, said, ‘...passing the Turing Test is not a sensible research and development goal for serious AI’”.

However, natural language processing is an extremely complicated endeavor, with a dauntingly high number of degrees of freedom, almost limitless combinations of units of meaning, and many obstacles to large scale knowledge representation (Saygin, Cicekli, & Akman, 2000). On the other hand, one can constrain the Turing Test to contain many fewer degrees of freedom. For example, in the current set of studies, we restricted the Turing test to either one degree of freedom in Experiment 1 (on or off switch) or to three

degrees of freedom in Experiment 2 (pitch, yaw, and roll). In this type of restricted test, a mimicking computer agent actually “passes” the nonverbal Turing Test, in the traditional sense in Experiment 1 and in a more liberal sense in Experiment 2. In many ways, an algorithm that employs simple mimicry is against the spirit of the theoretical foundation of the Turing Test, since it is not an “intelligent” algorithm that understands intentions, agency, and context. Nonetheless, in extremely simple instantiations of the test, mimicry appears to be an extremely successful substitute for artificial intelligence.

8.2. *Mimicry and theories of interactivity*

Algorithmic mimicry can be thought of as harnessing the natural human tendency to mimic the gestures and linguistic cues of people we interact with as a way to enhance the trust and rapport in an interaction (Chartrand & Bargh, 1999; Kendon, 1970). Kendon’s interactional synchrony and perhaps even behavioral contagion, the tendency for multiple members of groups to engage in simple behaviors such as yawns and laughter simultaneously (Neumann & Strack, 2000) in natural human behavior, are likely to be the observable effects of automatic mimicry. Past research has made it clear that automatic mimicry is pervasive in human interaction and functions as a regulator of trust between interactants. The current findings suggest that mimicry could easily be added to computer agents to improve the user’s experience unobtrusively – that is to say, without the user’s notice.

Current findings also tie into theories of interactivity. Conceptual definitions of interactivity typically emphasize three dimensions: technology, process, and user. Proponents of the technology dimension argue that interactivity is an affordance of technology (Steuer, 1992), while proponents of the process dimension argue that it is the mode of communication that influences interactivity – i.e. one-way versus two-way communication (Rafaeli, 1988). Finally, proponents of the user dimension contend that interactivity is a function of perception, and is ultimately in the mind of the participant (Kioussis, 2002; Wu, 1999). In all three views, feedback and contingency or expected contingency are the defining features of interactivity. Thus, agents employing automated mimicry may be more interactive than normal agents because they provide more contingent feedback to the user.

The Gricean Maxims of Conversation (Grice, 1975) provide another framework for understanding algorithmic mimicry. A mimicking agent fulfills the maxim of relevance – the expectation of appropriate non-verbal reactions after something is said or a non-verbal cue is given. Moreover, the mimicking agent also fulfills the maxim of manner – providing non-verbal cues that match the content of the conversation. In addition, the agent would also correctly reproduce the appropriate length and intensity of gestures and fulfill the maxim of quantity. However, the mimicking agent violates the maxim of quality, which in essence says “tell the truth”. This is because the mimicking agent is providing non-verbal cues designed to be manipulative and deceptive. The mimicking agent is essentially lying with non-verbal cues, and this violation of conversation maxims may be what lowers its perceived credibility and warmth when detected.

In terms of the maxim of relevance, the findings from the current studies question the nature of interactivity and contingency (see Burgoon et al., 2000 for a review of this concept). In Experiment 2, most would clearly indicate that agents in the mirror-mimic condition were in fact mimicking the participants. However, in the axis-switch condition,

it is not clear that the head movement behaviors, while perfectly contingent and highly interactive, was by definition a mimic. As behavioral contingencies become more abstracted in the use of digital media it becomes unclear, even when the contingencies remain perfect, what defines a mimic behavior. Consider a scenario in which modalities are switched – every time a user smiles, the agent's skin tone changes via a direct mathematical transform. This behavior is contingent, but is it a mimic? The levels at which embodied agents can abstract mimics using contingent behaviors is limitless.

8.3. Limitations, implications, and future directions

In the current research, we examined the ability of humans to detect mimicry behavior of embodied agents during interaction. However, we have implemented mimicry strategies in the most rudimentary manner – exact mimics in terms of scale (i.e. the size of the mimicked head movements compared to the actual head movements) and frequency (i.e., how often the head moved during mimicry compared to the physical movements) at differential delays. Furthermore, our tasks were extremely reduced to either one or three degrees of freedom. This is a poor approximation of how humans mimic one another in face-to-face interaction; consequently, our participants may have detected this inelegant mimic more often than they would have with a mimic that more closely modeled how humans interact normally, especially in Experiment 2. In future work, we plan to use more involved and generalizable mimic strategies, for example, ones that increase or decrease the rate of mimicry in response to certain actions by the user, or alternatively ones that mimic on a more probabilistic basis. Furthermore, we only utilized very simple mediated behaviors such as hitting a button or moving one's head. Given current abilities to track more complex nonverbal behaviors, such as face tracking (Hu et al., 2004) and eye gaze (Monk & Gale, 2002), we should soon be able to explore human interactions with agents that demonstrate realistic emotional mimicry.

In addition, future work should focus on locating the actual threshold or boundary that separates successful, implicit mimicry from the type of mimicry which is detected and consequently backfires. Given that the benefits of mimicry are extremely effective when implicit, but counterproductive when detected, understanding where this boundary lies has large theoretical and applied consequences. Along similar lines, exploring the difference between simple, linear combinations of multiple dimensions of mimicry (e.g., facial expressions, head movements, and other behaviors), and more elegant non-linear combinations of the same dimensions may help shed light on the location of this boundary.

Together, these two studies demonstrate the viability of using algorithmic mimics in computer agents. As multimodal input devices – such as voice recognition, expression recognition, or gesture recognition – increase in prevalence, the ability for computer agents to mimic human behaviors can become very sophisticated and layered. It also creates the potential for computer agents to mimic more complex human traits – such as speech patterns, personality, speech accent, expressiveness and so on. Research has shown that users with dominant personalities prefer computer agents that use preprogrammed dominant phrases (e.g. “You should definitely become a doctor”), while user with submissive personalities prefer computer agents that use preprogrammed submissive phrases (e.g. “Perhaps you should explore the medical profession”) (Moon & Nass, 1996; Reeves & Nass, 1996). Sophisticated algorithmic mimicry may provide one means of achieving this personality adaptation automatically without having to parse and understand personality.

For example, mimicry of verbosity, greeting phrases, speech patterns, expressiveness, and intonation may provide a very good approximation for apparent personality for computer agents. Given the evidence from previous work that mimicry causes social influence, as well as findings from the current work that demonstrate mimicry as a social influence strategy backfires when it is detected, interactants in computer-mediated discourse have a definite interest in detecting mimicry.

Acknowledgements

The authors thank Jakki Bailey, Claire Carlson, Alexandra Davies, and Jonathon Shih for their assistance in this research, Tonya Chartrand and Rick Van Baaren for helpful suggestions, and Alexandra Davies and Andrew Orin for comments on this paper.

Appendix

Agreement

- (1) I agree with the plan to implement ID cards.
- (2–4) I think the proposed ID cards are (valuable/workable/needed).

Agent Impression

- (1) The presenter was friendly.
- (2–12) The presenter was (likeable/honest/competent/warm/informed/credible/modest/approachable/in teresting/ trustworthy/sincere).

Open Ended Questions

Please list any thoughts you may have about the interaction with the presenter.

Was there anything unusual about this interaction?

Please write a few sentences about the presenter's LIP movements while speaking.

Please write a few sentences about the presenter's HEAD movements while speaking.

References

- Avouris, N. M., Dimitracopoulou, A., & Komis, V. (2003). On analysis of collaborative problem solving: an object-oriented approach. *Computers in Human Behavior*, *19*(2), 147–167.
- Bailenson, J. N., & Beall, A. C. (2006). Transformed social interaction: Exploring the digital plasticity of avatars. In R. Schroeder & A. Axelsson's (Eds.), *Avatars at work and play: Collaboration and interaction in shared virtual environments* (pp. 1–16). Springer-Verlag.
- Bailenson, J. N., Beall, A. C., Loomis, J., Blascovich, J., & Turk, M. (2004). Transformed social interaction: decoupling representation from behavior and form in collaborative virtual environments. *PRESENCE: Teleoperators and Virtual Environments*, *13*(4), 428–441.
- Bailenson, J. N., & Yee, N. (2005). Digital chameleons: automatic assimilation of nonverbal gestures in immersive virtual environments. *Psychological Science*, *16*, 814–819.
- Blascovich, J., Loomis, J., Beall, A., Swinth, K., Hoyt, C., & Bailenson, J. N. (2002). Immersive virtual environment technology as a methodological tool for social psychology. *Psychological Inquiry*, *13*, 103–124.
- Bourhis, R. Y., Giles, H., & Lambert, W. E. (1975). Social consequences of accommodating one's style of speech: a cross-national investigation. *International Journal of the Sociology of Language*, *6*(5), 5–71.
- Buller, D. B., & Burgoon, J. K. (1996). Interpersonal deception theory. *Communication Theory*, *6*, 203–242.
- Burgoon, J. K., Bonito, J. A., Bengtsson, B., Cederberg, C., Lundberg, M., & Allspach, L. (2000). Interactivity in human-computer interaction: a study of credibility, understanding, and influence. *Computers in Human Behaviour*, *16*, 553–574.

- Cappella, J., & Panalp, A. (1981). Talk and silence sequences in informal conversations. Interspeaker influence. *Human Communication Research*, 7, 117–132.
- Cassell, J. (2001). Representation and Intelligence in Embodied Conversational Agents. *AI Magazine*, 22(3), 67–83.
- Chartrand, T. L., & Bargh, J. A. (1999). The chameleon effect: the perception-behavior link and social interaction. *Journal of Personality and Social Psychology*, 76, 893–910.
- Cohen, E., & Cooper, R. (1986). Language and tourism. *Annals of Tourism Research*, 13, 533–563.
- Dovidio, J. F., Kawakami, K., & Beach, K. R. (2001). Implicit and explicit attitudes: examination of the relationship between measures of intergroup bias. In R. Brown & S. L. Gaertner (Eds.), *Blackwell handbook of social psychology. Intergroup relations* (Vol. 4, pp. 175–197). Oxford, UK: Blackwell.
- French, R. M. (2000). Peeking behind the screen: the unsuspected power of the standard Turing test. *Journal of Experimental and Theoretical Artificial Intelligence*, 12, 331–340.
- Fussell, S. R., Setlock, L. D., Yang, J., Ou, J., Mauer, E. M., & Kramer, A. (2004). Gestures over video streams to support remote collaboration on physical tasks. *Human-Computer Interaction*, 19, 273–309.
- Giles, H. (1980). Accommodation theory: some new directions. In S. de Silva (Ed.), *Aspects of linguistic behavior*. New York: University of York.
- Giles, H., Coupland, J., & Coupland, N. (1991). Accommodation theory: communication, context, and consequence. In H. Giles, J. Coupland, & N. Coupland (Eds.), *Contexts of accommodation. Developments in applied sociolinguistics*. Cambridge: Cambridge University Press.
- Giles, H. (1973). Communicative effectiveness as a function of accented speech. *Speech Monographs*, 40, 330–331.
- Graesser, A. C., Jackson, G. T., Mathews, E. C., Mitchell, H. H., Olney, A., Ventura, M., et al. (2003). Why/autotutor: a test of learning gains from a physics tutor with natural language dialog. In D. Hirsh & R. Alterman (Eds.), *Proceedings of the 25rd annual conference of the cognitive science society* (pp. 1–5). Boston, MA: Cognitive Science Society.
- Graf, P., & Schacter, D. L. (1985). Implicit and explicit memory for new associations in normal and amnesic patients. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11, 501–518.
- Grice, H. P. (1975). Logic and conversation. In P. Cole & J. Morgan (Eds.), *Syntax and semantics: Speech acts* (Vol. 3). New York: Academic.
- Guadagno, R. E., & Cialdini, R. B. (2002). Online persuasion: an examination of gender differences in computer-mediated interpersonal influence. *Group Dynamics: Theory Research and Practice. Special Issue on Internet Research*, 6, 38–51.
- Guadagno, R., Bailenson, J. N., Beall, A. C., Dimov, A., & Blascovich, J. (2005). Transformed social interaction and the cyranoid: the impact of non-verbal behavior on persuasion in an immersive virtual environment. In *Presentation to the conference on group processes in computer-supported interaction*, Miami, OH.
- Hu, C., Chang, Y., Feris, R. & Turk, M. (2004). Manifold based analysis of facial expression. In *IEEE workshop on face processing in video*, Washington, DC.
- Johnston, L. (2002). Behavioral mimicry and stigmatization. *Social Cognition*, 20, 18–34.
- Kahneman, D., & Tversky, A. (1973). On the psychology of prediction. *Psychological Review*, 80, 237–251.
- Kendon, A. (1970). Movement coordination in social interaction: some examples described. *Acta Psychologica*, 32(2), 101–125.
- Kiesler, S., Siegel, J., & McGuire, T. (1984). Social psychological effects of computer mediated communication. *American Psychologist*, 39, 1123–1134.
- Kiouis, S. (2002). Interactivity: a concept explication. *New Media & Society*, 4(3), 355–383.
- Konijn, E. A., & Hoorn, J. F. (2004). Reality-based genre preferences do not direct personal involvement. *Discourse Processes*, 38(2), 219–246.
- Kuiken, D., Phillips, L., Gregus, M., Miall, D. S., Verbitsky, M., & Tonkonogy, A. (2004). Locating self-modifying feelings within literary reading. *Discourse Processes*, 38(2), 267–286.
- LaFrance, M. (1979). Nonverbal synchrony and rapport: analysis by the cross-lag panel technique. *Social Psychology Quarterly*, 42, 66–70.
- Levelt, W. J. M., & Kelter, S. (1982). Surface form and memory in question answering. *Cognitive Psychology*, 14, 78–106.
- Loftus, E. (1996). *Eyewitness testimony*. Cambridge, MA: Harvard University Press.
- Monk, A. F., & Gale, C. (2002). A look is worth a thousand words: full gaze awareness in video-mediated conversation. *Discourse Processes*, 33, 257–278.
- Moon, Y., & Nass, C. I. (1996). Adaptive agents and personality change: complementarity versus similarity as forms of adaptation. In Bilger, R., Guest, S., Tauber, M.J. (Hrsg.), *Human factors in computing systems: Chi'96 electronic conference proceedings*.

- Morley, I., & Stephenson, G. (1969). Interpersonal and interparty exchange. A laboratory simulation of an industrial negotiation at the plant level. *British Journal of Psychology*, 60, 543–545.
- Natale, M. (1975). Convergence of mean vocal intensity in dyadic communications as a function of social desirability. *Journal of Personality and Social Psychology*, 32, 790–804.
- Neumann, R., & Strack, F. (2000). Mood contagion: the automatic transfer of mood between persons. *Journal of Personality & Social Psychology*, 79, 211–223.
- Pew Internet and American Life Project, P. I. A. L. (2004). *53 Million in US Use IM*, from <<http://www.emarketer.com/Article.aspx?1003031>>.
- Provine, R. (1992). Contagious laughter: laughter is sufficient stimulus for laughs and smiles. *Bulletin of the Psychonomic Society*, 30(1–4).
- Rafaeli, S. (1988). Interactivity: from new media to communication. In R. P. Hawkins, J. M. Wiemann, & S. Pingree (Eds.), *Advancing communication science: Merging mass and interpersonal processes* (pp. 110–134). Beverly Hills, CA: Sage.
- Reeves, B., & Nass, C. (1996). *The media equation: How people treat computers, television, and new media like real people and places*. Cambridge, MA: Cambridge University Press.
- Saygin, A. P., Cicekli, I., & Akman, V. (2000). Turing Test: 50 years later. *Minds and Machines*, 10(4), 463–518.
- Shamp, S. A. (1991). Mechanomorphism in perception of computer communication partners. *Computers in Human Behavior*, 7, 147–161.
- Short, J., Williams, E., & Christie, B. (1976). *The Social Psychology of Telecommunications*. London: Wiley Press.
- Steuer, J. (1992). Defining virtual reality: factors determining telepresence. *Journal of Communication*, 42(4), 73–93.
- Suzuki, N., Takeuchi, Y., Ishii, K., & Okada, M. (2003). Effects of echoic mimicry using hummed sounds on human–computer interaction. *Speech Communication*, 40, 559–573.
- Turing, A. M. (1950). Computing machinery and intelligence. *Mind*, 59, 433–460.
- Welkowitz, J., Feldstein, S., Finkelstein, M., & Aylesworth, L. (1972). Changes in vocal intensity as a function of interspeaker influence. *Perceptual and Motor Skills*, 35, 715–718.
- Wiener, M., & Mehrabian, A. (1968). *Language within language: Immediacy, a channel in verbal communication*. New York: Appleton-Century-Crofts.
- Whittaker, S. (2002). Theories and methods in mediated communication. In A. Graesser, M. Gernsbacher, & S. Goldman (Eds.), *The handbook of discourse processes* (pp. 243–286). NJ: Erlbaum.
- Wu, G. (1999). Perceived interactivity and attitude toward website. In M. S. Roberts (Ed.), *Proceedings from the 1999 American Academy of Advertising*. Gainesville, FL: University of Florida.
- Van Baaren, R. B., Holland, R. W., Steenaert, B., & van Knippenberg, A. (2003). Mimicry for money: behavioral consequences of imitation. *Journal of Experimental Social Psychology*, 39(4), 393–398.
- Van Baaren, R. B., Holland, R. W., Kawakami, K., & van Knippenberg, A. (2004). Mimicry and pro-social behavior. *Psychological Science*, 15, 71–74.
- Yee, N. (2006). The psychology of MMORPGs: Emotional investment, motivations, relationship formation, and problematic usage. In R. Schroeder & A. Axelsson (Eds.), *Avatars at work and play: Collaboration and interaction in shared virtual environments* (pp. 187–207). London: Springer-Verlag.
- Zuckerman, M., DePaulo, B. M., & Rosenthal, R. (1981). Verbal and non-verbal communication of deception. In L. Berkowitz (Ed.), *Advances in experimental and social psychology* (Vol. 14, pp. 1–59). New York: Academic Press.