Evaluating perceived trust from procedurally animated gaze

Aline Normoyle, Jeremy B. Badler, Teresa Fan, Norman I. Badler *
University of Pennsylvania
Vinicius J. Cassol, Soraia R. Musse †
Pontificia Universidade Catolica do Rio Grande do Sul

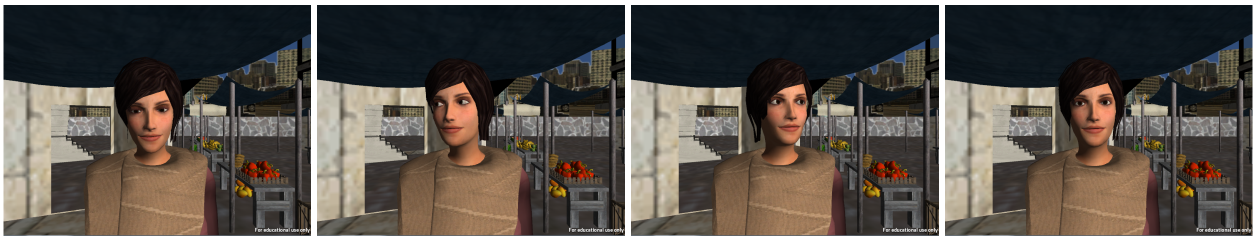

Figure 1: *Eye gaze model for signalling how much a non-player character (NPC) trusts the player.*

Abstract

Adventure role playing games (RPGs) provide players with increasingly expansive worlds, compelling storylines, and meaningful fictional character interactions. Despite the fast-growing richness of these worlds, the majority of interactions between the player and non-player characters (NPCs) still remain scripted. In this paper we propose using an NPC's animations to reflect how they feel towards the player and as a proof of concept, investigate the potential for a straightforward gaze model to convey trust. Through two perceptual experiments, we find that viewers can distinguish between high and low trust animations, that viewers associate the gaze differences specifically with trust and not with an unrelated attitude (aggression), and that the effect can hold for different facial expressions and scene contexts, even when viewed by participants for a short (five second) clip length. With an additional experiment, we explore the extent that trust is uniquely conveyed over other attitudes associated with gaze, such as interest, unfriendliness, and admiration.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Animation

Keywords: facial animation, procedural eye gaze, perception, trust

*alinen,jbadler,terfan,badler@seas.upenn.edu
†vinicius.cassol,soraia.musse@pucrs.br

Motion in Games 2013, November 7 – 9, 2013, Dublin, Ireland.

1 Introduction

Adventure role playing games (RPGs) provide players with increasingly expansive worlds, compelling storylines, and meaningful fictional character interactions. Although such games give players increasing agency to directly influence the virtual world through their game actions, the majority of interactions between the player and non-player characters (NPCs) still remain static. For example, the player will walk repeatedly into the same inn, talk to the same vendors and receive the same dialog choices every time. Most interactions are not yet personalized according to the player's appearance, race, class, or quest results [Vanhatupa 2011]. Exceptions to this rule tend to be labor-intensive since they usually rely on scripting, and thus require additional design and QA effort to anticipate potential player choices. Personalizing game-play via quests and storyline is particularly challenging in online RPGs because quest objectives must periodically "respawn" so that other players can play them. However, personalizing non-combat NPC interactions, which can be directed towards a specific player, do not suffer from this limitation. Thus, we propose using the animation of characters to reflect their opinions towards the player, such that no additional scripting is needed.

There are numerous aspects of interactions that are amenable to personalization. In this paper, we evaluate the potential for a straightforward gaze model to convey an NPC attitude toward a player, specifically trust. We focus our proof of concept on conveying trust because it is a fundamental aspect of inter-personal relationships which manifests in numerous non-verbal behaviors [Vrij and Mann 2004; Frith and Frith 2007; Willis et al. 2011]. Trust intuitively links to the concept of reputation, which is a popular game mechanic in many RPGs, including Star Wars: The Old Republic, World of Warcraft, and EVE Online. In these games, as a player's reputation with an NPC's group increases, he gains access to better gear, better missions, and better rewards; however, these NPCs provide no behavioral feedback to show that their *attitude* towards the player has changed. The ability to influence the behaviors of group NPCs could potentially be additionally rewarding and motivating for players who enjoy role-playing and story immersion [Yee 2006; Thue et al. 2011]. Lastly, because trust is conveyed effectively with the face, it is appropriate for most game environments, which zoom or center into the faces of NPCs when the player speaks to them.

In our approach, we implement a naturalistic gaze model supporting eye saccades, lid movement, and blinks. The rates and durations of eye and lid movements are determined probabilistically based on rates from existing physiology literature. The gaze model is parameterized in terms of trust, which probabilistically modulates the proportion of time the NPC looks at either the player or at a peripheral target. Although some work with virtual characters has already investigated whether the player trusts the character, we are interested in the converse question: how much does the player think the character trusts them? With this goal in mind, to try to understand the capabilities and limitations of the model we ask the following questions:

- Can viewers distinguish high-trust from low-trust animations?
- Do viewers associate the animations with trust rather than an unrelated attitude towards the player?
- Does the expression of the character affect perceived trust?
- Does the scene context, as conveyed through text character dialog, affect perceived trust?
- Is it possible to convey something meaningful in a short exposure time, such as those typically found in game environments?
- How consistent are attributions of trust across individuals?

In light of these questions, we evaluated the NPC gaze model using three perceptual experiments. In the first, we tested the model using an animated character with a friendly neutral face. In the second, we repeated the experiment with a grumpy, or unfriendly, neutral face. In both experiments, we exposed viewers to a series of five second clips, prefaced with a text dialog, in which participants had to indicate either the perceived level of trust or the perceived level of aggression, an attitude which can also be attributed to head position and gaze [Grant 1969; Ellsworth et al. 1972; Adams Jr and Kleck 2005]. In the third experiment, we again showed participants five second clips, but asked participants to rate the perceived level of trust, aggression, unfriendliness, admiration, or interest conveyed by the character towards the player. We found that:

- Yes, viewers can distinguish between high and low trust animations.
- Viewers associate the gaze differences specifically with trust and not with an unrelated attitude (aggression). However, this effect is not sufficiently powerful to uniquely convey different, related attributes, such as admiration or interest .
- The effect can hold for different facial expressions.
- The effect is independent of scene context.
- Viewers can make judgements of trust using a clip length as short as 5 seconds.
- It's not just a population averaging effect; it is consistent across participants, holding up for the majority of subjects even when they are considered individually.

In the following sections, we will outline related work and background from the psychology literature. We will then describe our gaze and eye model and detail the execution and analysis of our perceptual experiments.

2 Related Work

Numerous systems have been proposed for animating speech, expressions, and eye movements [Lee et al. 2002; Masuko and Hoshino 2007; Wang et al. 2007; Deng and Neumann 2008; Queiroz et al. 2008; Ahn et al. 2012]. For a survey of techniques suitable for games, see [King 2008]. In particular, approaches for animating gaze and eyes tend to be probabilistic, corresponding to either parameterized models (usually constructed from empirical data, such as [Lee et al. 2002; Itti et al. 2003; Masuko and Hoshino 2007]) or models constructed automatically using statistical learning techniques, such as [Deng et al. 2005; Ma and Deng 2009; Le et al. 2012]. In this work, we use parameterized models to animate gaze, eye movements, and blinks, based on existing research (see section 4).

Gaze and eye movements are essential for non-verbal communication. Early work in embodied conversational agents [Cassell and Thorisson 1999] established the importance of envelope feedback (gaze, beat gestures, and head movements) for dialog with humans and found that it was more crucial than expressions. Similarly, early work in pedagogical agents [Johnson et al. 2000] established the ability of gaze to guide student attention and of nods and head shaking to provide non-obtrusive student feedback. Other work has looked at how non-verbal cues can indicate extroversion and introversion [Isbister and Nass 2000], display emotions [Lance and Marsella 2008], indicate arousal and dominance [Cig et al. 2010], indicate attention and distraction [Gu et al. 2007], facilitate establishing common ground [Nakano et al. 2003], show puzzlement and confusion [Stone and Oh 2008], and build rapport [Gratch et al. 2007]. [Kulms et al. 2011] looked at whether gaze behaviors were perceived differently on either male or female characters and determined that gaze aversion was viewed equally negatively for both.

Most work in embodied conversational agents are interested in real-time multi-modal interactions between virtual characters with humans (for example, the use of speech and face recognition to detect when the user has stopped speaking). Although such interfaces are not part of current games, existing work has begun to explore this potential [Reale et al. 2011]. However, we are interested in the more modest of goal of integrating a richer repertoire of non-verbal cues into NPC characters, for the purposes of reflecting the game attributes of the player.

Gaze behavior has also been shown to influence the outcome of human-avatar interactions. For example, [Bailenson et al. 2005] created a networked collaborative environment in which gaze was manipulated so that all participants thought the speaker was focused exclusively on them. This manipulation increased the persuasiveness of the speaker's message for women and increased the recall of the speaker's message in men. [Bailenson and Yee 2005] showed that mimicry in virtual agents induced positive social effects similarly to mimicry in real life, increasing character persuasion and positive ratings. [Steptoe et al. 2010] studied the impact of the relative eye motion in an avatar mediated communication and found that the addition of eye movement increased participant's ability to detect truth and deception while interacting with avatars. [Krämer et al. 2013] found that increases in character smiling increased smiling in viewers (but did not change the viewer's opinion of the avatar). [Bee et al. 2009] studied whether flirting could improve first impressions between a user and a virtual character.

Two additional studies included a mechanism for trust, but neither used eye gaze for its implementation. Martinez *et al* [2008] specifically modeled the level of trust that an agent had in a human collaborator, using verbal responses to reflect its state. [Colburn et al. 2000] built a predictive model of trust based on voice for the purposes of guiding user interface interactions based on the user's mental state.

3 Background: Eye gaze and trust

How can an agent nonverbally communicate its level of trust towards a human user? One of the most important channels of nonverbal communication is eye gaze, which is used by humans as well as animals to convey a wide rage of information [Kleinke

1986; Emery 2000]. The timing and direction of gaze can indicate subtle characteristics and states such as dominance or submissiveness, where a person's attention is at the moment, flirtation, interest in beginning a conversation (or desire to avoid one), an invitation for one's conversation partner to take a turn in the dialogue, active listening, and pondering a point [Isbister 2006]. [Willis et al. 2011] examined how gaze direction along with facial expression impact the approachability and trustworthiness of faces. They found that direct gaze significantly increased perceived trustworthiness for happy and neutral faces, but decreased it for angry ones. Another study [Wyland and Forgas 2010] also reported that direct gaze increased trustworthiness, regardless of the mood of the subject. Faces that consistently looked at invalid objects during a classification task were unsurprisingly rated as untrustworthy [Bayliss and Tipper 2006].

There is an important distinction between conveying trustworthiness (how much do I trust someone) and trust (how much does someone trust me). The latter has been less well studied than the former, but there is still substantial evidence that it is tightly linked to gaze. For example, human observers tend to prefer items that a trusted individual looks at, as opposed to an untrusted one [King et al. 2011]. They also make shorter duration eye fixations while viewing untrusted websites [Leichtenstern et al. 2011]. Individuals cooperating in a game look at each other approximately twice as long as individuals engaged in competition [Foddy 1978]. Importantly, when humans engage in social interactions they tend to mimic one another's gestures [Chartrand and Bargh 1999], so it is reasonable to assume that eye movements conveying trust and trustworthiness are similar. Note that trust is not the same as rapport, as it implies some degree of vulnerability in the trusting person [Pearce 1974].

Based on the aforementioned studies, we derived the following model whereby a virtual agent could express high versus low trust toward a human player. The high trust state is characterized by a higher percentage of eye contact (looking at the player) when the player is looking at the agent. Low trust is indicated by the opposite, i.e. more frequent averted gaze when the player looks at the agent. The potential confounding factor of facial expression [Oosterhof and Todorov 2008] was held constant over an experimental session, and no audio was used to avoid potential speech cues [Vrij and Mann 2004]. The trust model was evaluated by asking participants to rate how much they thought the agent trusted them, using 10 different text dialogs. On half of the trials a control question unrelated to trust (on aggression) was asked instead, to provide an additional reference for comparison.

4 Gaze and Blink Model

To evaluate the role of gaze for conveying trust, we built a plausible naturalistic gaze and eye movement model based on existing physiology literature.

Blinks and eye lid movements

Eyelids have been shown to move synchronously with the vertical movements of the eyes, called lid saccades [Evinger et al. 1991]. Such movements tend to follow a linear relationship with vertical eye rotation. We approximate this effect with a simple ratio, configured for both upward and downward eye movements.

$$e^t = e_{\max}\frac{\theta^t}{\theta_{\max}}$$

where e^t is the lid displacement for this time step, $e_{\max}$ is the maximum lid displacement, θ^t is the current vertical eye rotation, and

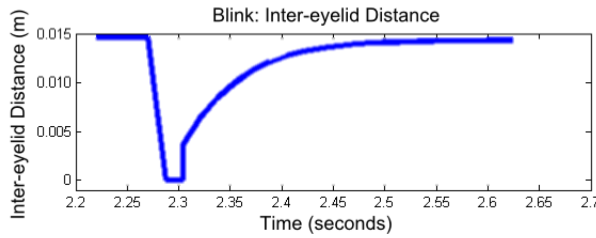

Figure 2: *Blinks. This curve shows inter-eyelid distance as the lid quickly closes in the down phase and then asymptotically lifts back to its original height.*

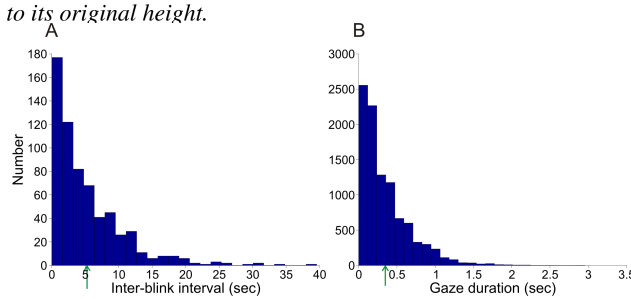

Figure 3: *Histograms of gaze model performance for all trials of the first experiment. Blink and gaze durations are exponentially distributed. Green arrows indicate the means. A: durations between blinks (theoretical mean is 11.6 and simulated mean 11.37). B: gaze durations (or equivalently, inter-saccadic intervals). The simulated mean is 348 msecs.*

$\theta_{\max}$ is the maximum vertical eye rotation. The lid displacement is then added to the base eyelid height.

Blink dynamics have been shown to be asymmetric: eyelids typically close very quickly and then asymptotically return to their original position [Evinger et al. 1991; Trutoiu et al. 2011]. We approximate these dynamics with a fast linear lid down-phase, but allow the up-phase to be controlled by the lid saccade (Figure 2). This approach is supported by [Evinger et al. 1991], who report that the kinematics of upward saccades and upward blinks are virtually identical.

In normal individuals, blink rates can vary due to mental state, emotional state, or task [Trutoiu et al. 2011]. We choose a rate of 11.6 blinks, reported to be a typical spontaneous rate during conversations under normal conditions [Itti et al. 2003; Doughty 2001]. The interval between blinks is then modeled with an exponential distribution with mean 11.6 per minute. A histogram of all blink intervals simulated in the first experiment is shown in Figure 3, panel A. The simulated mean rates for the first and second experiments were 11.37/min and 11.57/min respectively.

Gaze and saccades

Gaze shifts consist of moving the head and the eyes in the same direction [Guitton and Volle 1987; Itti et al. 2003]. For small gaze shifts, predominantly the eyes move; for larger gaze shifts, both the head and eyes move [Tomlinson and Bahra 1986]. We implement this effect with a simple user-defined threshold of ten degrees. Head movements are computed using inverse kinematics of the head and neck joints.

Saccades are implemented as small gaze shifts between close gaze targets. When our character is looking at a face, we define these gaze targets as the eyes, nose, and mouth [Buchan et al. 2007]. When looking away from the face, the character looks at precomputed offset locations from the central point of interest (see Figure 4). For example, if the character looks at a landmark to the left of the player's head, the saccades will occur at four pre-computed positions relative to the landmark. Because saccades are extremely

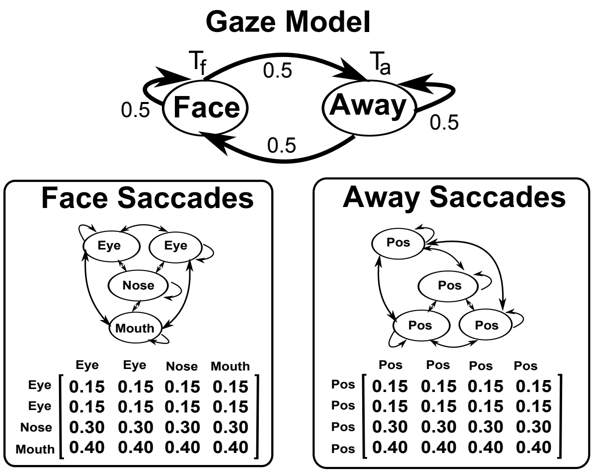

Figure 4: *Gaze state and saccade transition graphs. The gaze model consists of two states. The durations in each state, T_a, T_f are determined by trust level. Saccades are implemented with probabilistic transitions between positions. For the face, these positions are facial features; for objects in the world, they are offsets from the object's center. When a saccade completes, we transition to a new position according to the probabilities shown in the corresponding matrix. Each column corresponds to the transition probabilities from each location. For example, for face saccades, the first column shows the transition probabilities from an eye to all other face locations.*

fast (with accelerations as high as 30,000 deg/s and peak velocities of 400-600 deg/sec [Becker et al. 1989]), we do not model the velocity curves of their movements, but instead simulate them by directly rotating the eyes to a new position.

Between gaze shifts, the character's eyes are fixated on a single location. The duration of fixations follows an exponential distribution [Harris et al. 1988]. For the face, we configure the exponential model so that the proportion of fixations across facial features is slightly higher for the mouth, consistent with [Buchan et al. 2007] who showed that people tend to fixate there slightly more often when focusing on speech. This effect is achieved by making the probability of transitioning to the mouth higher than other face landmarks (Figure 4). Based on the same study, we also implement a mean fixation duration of 400 ms. In our implementation, when saccades are combined with gaze movements, the simulated mean rate will be faster. Our measured fixation times for experiment one are shown in Figure 3, panel B. Experiment two has a similar distribution; the means are 348 and 349 msec respectively. Note that these are slightly faster than the target of 400 ms, but still within the range of values recorded by [Buchan et al. 2007].

The trust model controls the proportion of time in which the character looks towards the viewer or gazes away (motivated by [Foddy 1978]). This is implemented probabilistically using an exponential model which transitions evenly between the two states (looking towards or away, see Figure 4). The level of trust then modulates the amount of time spent in each state. We allow trust to vary from -1 (low trust) to 1 (high trust). Specifically, if T_f and T_a are the durations in each state, let $D = T_f + T_a$. When trust is zero, we define $T_f = T_a$, e.g. the character spends an equal amount of time looking at the face or away. When trust is negative, $T_a = \alpha D$ and $T_f = (1 - \alpha)D$, with $0.5 < \alpha < 0.75$. When trust is positive, $T_a = (1 - \alpha)D$ and $T_f = \alpha D$, with $0.5 < \alpha < 0.75$. We clamp α at 0.75 so that head fixations do not become overly short at low and high trust states.

The distribution of gaze locations for a sample block of forty 5-second trials is plotted in Figure 5. The widths of the circles indicate the relative amount of time that the agent's gaze was fixated at a particular location. When the trust parameter was set to high (Panel A), there were more fixations on the facial features of the player: 64% in this example (inset). By contrast, a low trust parameter setting (Panel B) yielded more fixations on the peripheral targets, with player features comprising only 40%.

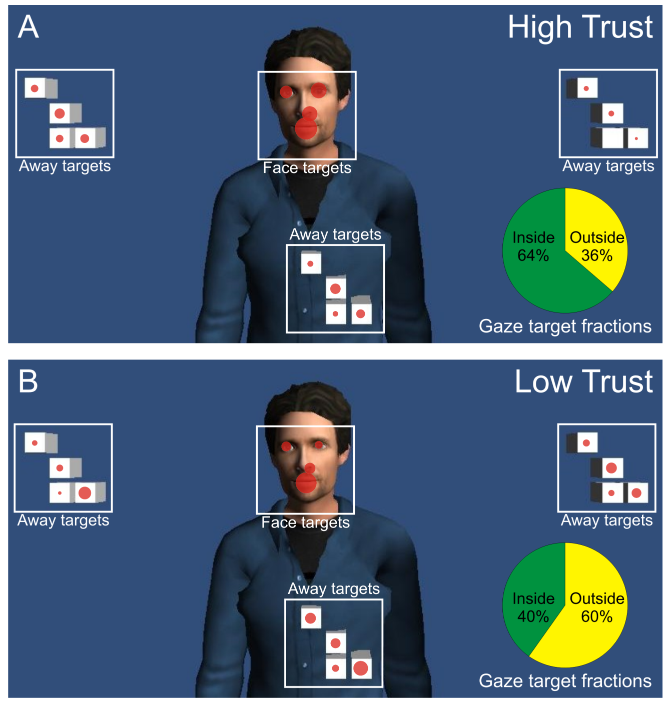

Figure 5: *Gaze locations for high trust and low trust trials for a single session (40 trials). The human player is shown from the NPC's perspective. Circle diameters indicate the relative amount of time that the NPC's gaze was fixated on a particular target. Total gaze fractions are indicated in the insets; inside targets are facial features and outside targets are the peripheral boxes. A: a high trust parameter yields more NPC fixations on the face of the player. B: a low trust parameter yields more peripheral fixations.*

5 Evaluation

The fundamental goal of the project was to convey to users how much an agent trusted them. In turn, trust level was manipulated by a single parameter within an eye gaze model. The most straightforward method of assessing whether the goal was successful was to simply show participants an agent whose gaze was driven by the model, then asking them to rate how much the agent seemed to trust them. We only assessed two trust levels, high and low, due to the constraints of the experimental environment: multiple factors had to be tested and repeated measurements of subject ratings were needed to obtain reliable results. For the same reason, exposure time to each condition was limited to five seconds.

Two important secondary factors to evaluate were robustness and specificity. In order to be robust, the players' assessment of agent trust had to be consistent across different scenario contexts (marketplace, town hall, etc.) and resistant to limited changes in facial expression. To manipulate context, we included 10 different text descriptions as an experimental factor. To test facial expression, we ran two separate versions of the experiment: one with a slightly happy facial expression and the other with a slightly negative facial expression (see Figure 7). In order to be specific, changes in the trust parameter needed to elicit changes in perceived trust and not

Figure 6: *Scene.*

(a) (b)

Figure 7: *Expressions used for the two scenarios.*

in another perceived trait. We tested specificity by asking participants to also assess agent aggressiveness, another emotion which can be communicated by gaze [Grant 1969; Ellsworth et al. 1972; Adams Jr and Kleck 2005]

5.1 Stimuli Creation

We use two fully articulated characters, one male and one female for this experiment. Only the female model is seen by participants, whereas the male model is used to represent the player.

We set up two scenarios intended to represent a player/NPC conversation in a virtual village (Figure 6) using the Unity game engine [unity3d.com]. In the first scenario, the character had a neutral friendly expression (Figure 7(a)). In the second scenario, the character had a neutral grumpy expression (Figure 7(b)). In both scenarios, we set up gaze targets *a priori*, as we did not implement an attention model. One look target was centered at the player character's face, with saccade points at the eyes, nose, and mouth, and three look targets were placed left, right, and below the character's face, with four saccade points specified around each. The scene camera was positioned such that the NPC character was centered in the view, level with the player's head.

This scene was incorporated into a randomized survey (Figure 8).

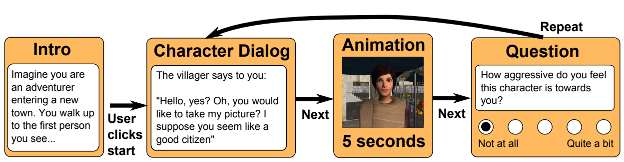

Figure 8: *Experiment Method. The introduction screen asks for gender and age information and gives context to the upcoming scenes. Each subsequent question shows one of 40 randomized trials. For each trial, we first show character dialog from the NPC in text, then a five second view of the NPC character, animated using our model, and finally a question regarding either aggression or trust rated on a five point scale.*

An introduction screen asked for the participant's gender and age followed by the description: "You walk up to the first townsperson you see and start a conversation. The next several scenes will show a series of potential reactions to you from the NPC. We would like you to rate what you believe the NPC's feelings are towards you." When the participant was ready, he or she clicked a button labeled "Begin". For each question, the participant was shown a short dialog from the NPC character with text, for example, "You look like you could use new equipment. I always reserve my best wares for dependable people." When ready, the participant clicked the Next button and he was shown a 5 second clip of our character, with the gaze and eye movements animated as described in section 4. We then asked for the participant's impression of how the NPC felt about the player using one of two probe questions: "How much do you feel this character trusts you?" or "How aggressive do you feel this character is towards you?". The participant responded on a scale of one to five and proceeded to the next trial.

5.2 Method

For each trial, the utterance, condition, and trust level (either high trust or low trust) was randomized. Thus each participant observed a total of 10 x 2 x 2 = 40 trials. Since gaze was determined probabilistically, the ratio of gaze towards versus away from the participant was also recorded.

Statistical analysis was performed using MATLAB [mathworks.com]. The two facial expression scenarios were analyzed separately. For each probe question, a planned, paired t-test was used to test whether responses in the low and high trust conditions were significantly different. A 2-way ANOVA with interactions was also performed on subject responses, using trust level and scene context as factors. Finally, the relation between response and gaze ratio for each trial was analyzed using linear regression.

5.3 Results

A total of 35 people participated in the experiment. Seventeen subjects (7 female, age range 21-53) viewed scenario one and the neutral-friendly face; 18 different subjects (5 female, age range 18-45) viewed scenario two and the neutral-grumpy face.

Figure 9 shows that the changes controlled by the trust parameter were indeed perceptible to observers. The pairs of bars show the average participant response for low (red) and high (blue) trust parameter settings. For both scenarios, subjects reported that the NPC trusted them more when the trust parameter was high (paired t-test; $p < 0.002$). Importantly, there was also evidence that the effect was specific: for the control question, where subjects had to rate how aggressive the NPC felt towards them, there was no effect of the trust parameter (paired t-test; $p > 0.5$).

In addition to the trust parameter and the probe question, the short utterance text used to set the scene for each trial was parametrically varied. Aside from making the experiment less tedious, the different texts allowed the assessment of whether the ratings were dependent on a specific scenario, or whether they were more general (i.e., dependent only on gaze). For both scenarios there was no significant marginal effect of utterance (2-way ANOVA; $p > 0.3$ for all), nor any interaction with trust parameter value or probe question (2-way ANOVA; $p > 0.1$ for all). The results indicate that trust rating was not sensitive to scene context.

From the perspective of the subjects, different values of the trust parameter were ultimately manifested as different gaze patterns produced by the agent. Specifically, as programmed, a high setting yielded more direct gazes and a low setting yielded more averted

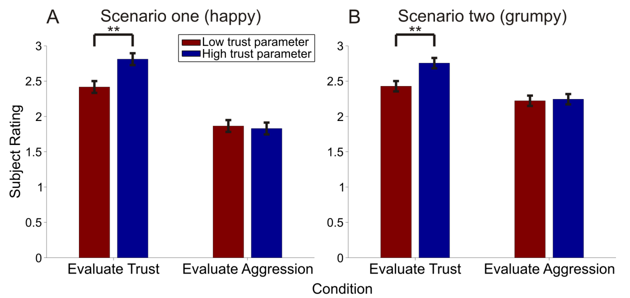

Figure 9: *Average ratings for different trust parameter settings and probe questions. When NPC trust toward the player was probed, the response was significantly different according to the trust parameter setting (left bars). By contrast, when NPC aggressiveness toward the player was probed, the value of the trust parameter had no effect (right bars). Error bars indicate standard error of the mean over all subjects. A: scenario one. B: scenario two.*

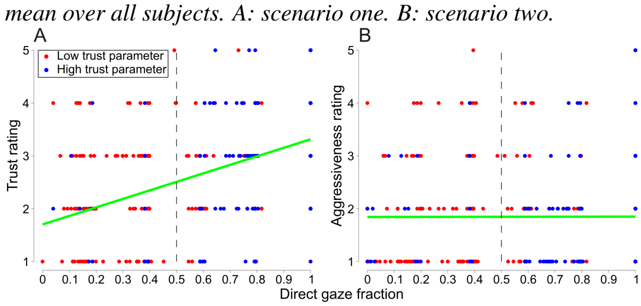

Figure 10: *Subject response versus direct gaze fraction for all trials of scenario one. Data points are color-coded based on whether they come from low or high trust parameter trials. Regression lines are plotted in green. A: probe question on NPC trust. Subject ratings were significantly correlated with the fraction of time the NPC gazed toward them, as shown by the regression line with positive slope. B: probe question on NPC aggressiveness. There was no correlation between rating and gaze fraction, as the regression line did not have a slope significantly different from zero.*

gazes. However, since the gaze model was probabilistic, not all trials with the same parameter setting yielded identical gaze fractions. To more precisely determine how actual gaze fraction influenced subject responses, the correlation between the two was calculated. Figure 10 plots subject ratings versus actual gaze fraction, for every trial in the scenario one experiment. Panel A shows there is a significant correlation for trust rating (green line; slope = 1.61, $p < 0.0001$). The result for scenario two was similar (not shown; slope = 1.20, $p < 0.0001$). In other words, subjects seemed to be aware of how much time that the avatar spent gazing at them during a trial, and it influenced their trust rating accordingly. Once again, it is apparent that the effect was specific to trust: when aggressiveness rating was plotted versus gaze fraction (Panel B) there was no correlation (scenario one: slope = 0.01, $p > 0.5$; scenario two not shown: slope = 0.28, $p > 0.1$).

Although there was a clear tendency for subjects as a whole to associate direct eye gaze with perceived trust, it is important to consider how reliable the effect was for individuals. Thus we performed the same regression analysis on each subject individually. For scenario one, ten out of 17 subjects showed a significant positive correlation between trust rating and gaze fraction, and six additional subjects had a positive though nonsignificant correlation. By contrast, only three subjects had significant correlations for the aggressiveness rating, and slopes were evenly divided between positive and negative. The results for scenario two showed a similar distribution: seven out of 18 subjects showed a significant positive correlation and seven more a nonsignificant positive correlation with trust. For aggressiveness, only three subjects had significant correlations

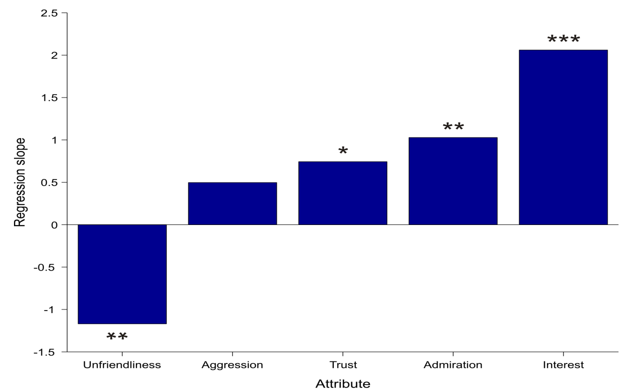

Figure 11: *Subject population regression slopes for the five attributes tested in the third experiment. The slopes measure the tendency for longer intervals of direct gaze to correlate with higher ratings of each particular attribute. Asterisks indicate degrees of significance: <0.01, <0.001 or <0.0001.*

and slopes were close to even (ten positive, eight negative). Overall, the results suggest that the perception of trust exists on more than just a population level: the majority of individual subjects are able to recognize the difference between simulated high-trust and low-trust gaze patterns, despite the random nature of the model and limited exposure times.

5.4 Limitations of the model

The simple direct/averted gaze model was successfully able to convey the impression of trust, as distinguishable from aggressiveness, in our limited experimental setup. However, we also wanted to probe the limitations of the model and identify specific weaknesses. To this end, we designed a third experiment that simultaneously tested three other attitudes besides trust and aggression: unfriendliness, interest and admiration. Similarly to our previous experiments, we control for facial expression. Since expression had no effect on trust perception, we arbitrarily chose the neutral grumpy expression. The experiment format was largely analogous to the first two: subjects viewed a brief introductory text to set the scene, then viewed 40 consecutive trials of NPC gaze, again five seconds each. At the end of each trial, the subjects had to answer the question "How *x* do you feel the NPC was toward you?" where *x* was a variant of the five previous attributes. A total of 21 people participated in the experiment. The results of the regression analysis (c.f. Figure 10) are shown in Figure 11. Consistent with the first two experiments, trust had a significant positive slope and the aggression slope did not differ significantly from zero. However, admiration and interest both had significant positive slopes as well, indicating that the gaze model alone was not sufficient to differentiate them from trust. Unfriendliness had a significant negative slope, indicating that its inverse (friendless) is also a likely confound. We can thus conclude that although the simple gaze model can convey information about a character attribute, it is not sufficiently powerful to uniquely convey different, related attributes.

6 Discussion

The trust model implemented here was greatly simplified compared to real human behavior. In real life, facial expressions [Oosterhof and Todorov 2008], mutual attention to objects [Bayliss and Tipper 2006] and verbal information [Pearce 1974] play extremely important roles in conveying the level of trust of one human in another. By contrast, the avatar's display of trust was essentially distilled to a single factor, namely the proportion of time they spent looking at the player. This was done for several reasons. First and foremost, we wanted to establish a minimal set of conditions by which the attitude of a virtual character could be effectively transmitted to a human player. Eye movements are easier to parameterize and model than other transmission modalities (e.g., facial expression), both conceptually and computationally. After showing eye movements to be sufficient, additional features can be introduced to increase the robustness of possible transmitted attitudes. Here we strived for a thorough evaluation of the capabilities of eye movements for conveying trust, leaving other modalities outside of the current scope.

Despite the impoverished representation of trust, the results (see Figure 9) indicated that the players were still able to perceive the difference between a high-trust and low-trust state. The perception is unlikely to be a simple artifact of correlating gaze to rating: for the control data, an analogous question (how aggressive is the NPC towards me) did not yield the same result. We thus conclude that our relatively simple eye movement model can be a suitable *basis* for conveying attitude. Nonetheless, as a simple model it has some drawbacks. First, as a probabilistic model it can be deceiving: a low trust parameter can still infrequently yield a high proportion of direct gaze, and vice-versa (see Figure 10). Adding additional information, such as a facial expression that is also modulated by the trust parameter, would help mitigate this effect. A second drawback is specificity: while the direct/averted gaze state might be sufficient to differentiate trust from aggression, it is not enough to differentiate trust from other attitudes such as interest or admiration. Once again, adding a repertoire of facial expressions corresponding to different attitudes is necessary to increase the power of the model. Additional enhancements of the model include implementing mutual attention, where the NPC is likely to look towards objects the player is looking at and implementing body language where the NPC turns towards or away from the player. Customizing dialog is also an option, though it requires additional scripting.

Although we can say that viewers correlated the eye gaze with higher trust when asked directly in an experimental setting, further experiments are needed to establish how noticeable the effect is to players immersed in a game. Nonetheless, this sort of approach could add to the player's agency and feeling of accomplishment: as the player progresses in the game, he can readily see the impact of his decisions on the virtual populace without the need for extensive developer involvement.

Acknowledgements

This work was partially supported by NSF grant CNS-0855192 and the Brazilian Research Agencies CAPES and CNPq. We wish to thank our participants and the anonymous reviewers for their comments.

References

ADAMS JR, R. B., AND KLECK, R. E. 2005. Effects of direct and averted gaze on the perception of facially communicated emotion. *Emotion 5*, 3.

AHN, J., GOBRON, S., THALMANN, D., AND BOULIC, R. 2012. Conveying real-time ambivalent feelings through asymmetric facial expressions. In *Motion in Games*. Springer.

BAILENSON, J., AND YEE, N. 2005. Digital chameleons: automatic assimilation of nonverbal gestures in immersive virtual environments. *Psychol Sci 16*, 814–9.

BAILENSON, J. N., BEALL, A. C., LOOMIS, J., BLASCOVICH, J., AND TURK, M. 2005. Transformed social interaction, augmented gaze, and social influence in immersive virtual environments. *Human Communication Research 31*, 511–537.

BAYLISS, A. P., AND TIPPER, S. P. 2006. Predictive gaze cues and personality judgments should eye trust you? *Psychological Science 17*, 514–520.

BECKER, W., ET AL. 1989. The neurobiology of saccadic eye movements. metrics. *Reviews of oculomotor research 3*, 13.

BEE, N., ANDR, E., AND TOBER, S. 2009. Breaking the ice in human-agent communication: Eye-gaze based initiation of contact with an embodied conversational agent. In *Intelligent Virtual Agents*. Springer Berlin Heidelberg.

BUCHAN, J., PAR, M., AND MUNHALL, K. 2007. Spatial statistics of gaze fixations during dynamic face processing. *Social Neuroscience 2(1)*, 1–13.

CASSELL, J., AND THORISSON, K. R. 1999. The power of a nod and a glance: Envelope vs. emotional feedback in animated conversational agents. *Applied Artificial Intelligence 13*, 519–538.

CHARTRAND, T. L., AND BARGH, J. A. 1999. The chameleon effect: The perception-behavior link and social interaction. *Journal of Personality and Social Psychology 76*, 893–910.

CIG, C., KASAP, Z., EGGES, A., AND MAGNENAT-THALMANN, N. 2010. Realistic emotional gaze and head behavior generation based on arousal and dominance factors. In *Motion in Games*, MIG'10.

COLBURN, A., COHEN, M. F., AND DRUCKER, S. 2000. The role of eye gaze in avatar mediated conversational interfaces. *Sketches and Applications, SIGGRAPH'00*.

DENG, Z., AND NEUMANN, U. 2008. *Data-driven 3D facial animation*. Springerverlag London Limited.

DENG, Z., LEWIS, J., AND NEUMANN, U. 2005. Automated eye motion using texture synthesis. *Computer Graphics and Applications, IEEE 25*, 24–30.

DOUGHTY, M. J. 2001. Consideration of three types of spontaneous eyeblink activity in normal humans: during reading and video display terminal use, in primary gaze, and while in conversation. *Optometry & Vision Science 78*, 712–725.

ELLSWORTH, P. C., CARLSMITH, J. M., AND HENSON, A. 1972. The stare as a stimulus to flight in human subjects: a series of field experiments. *Journal of personality and social psychology 21*, 302.

EMERY, N. 2000. The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience & Biobehavioral Reviews 24*, 581–604.

EVINGER, C., MANNING, K. A., AND SIBONY, P. A. 1991. Eyelid movements: Mechanisms and normal data. *Investigative Ophthalmology & Visual Science 32 (2)*, 387–400.

FODDY, M. 1978. Patterns of gaze in cooperative and competitive negotiation. *Human Relations 31(11)*, 925–938.

FRITH, C. D., AND FRITH, U. 2007. Social cognition in humans. *Current Biology 17*, R724–R732.

GRANT, E. C. 1969. Human facial expression. *Man 4*, 525–692.

GRATCH, J., WANG, N., GERTEN, J., FAST, E., AND DUFFY, R. 2007. Creating rapport with virtual agents. In *Intelligent Virtual Agents*.

GU, E., LEE, S., BADLER, J., AND BADLER, N. 2007. Eye movements, saccades, and multiparty conversations. In *Data-Driven 3D Facial Animation*. Springer London.

GUITTON, D., AND VOLLE, M. 1987. Gaze control in humans: eye-head coordination during orienting movements to targets within and beyond the oculomotor range. *Journal of neurophysiology 58*, 427–459.

HARRIS, C. M., HAINLINE, L., ABRAMOV, I., LEMERISE, E., AND CAMENZULI, C. 1988. The distribution of fixation durations in infants and naive adults. *Vision Research 28*, 419–432.

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*, 251 – 267.

ISBISTER, K. 2006. *Better Game Characters by Design A Psychological Approach*. Elsevier and Morgan Kaufmann.

ITTI, L., DHAVALE, N., AND PIGHIN, F. 2003. Realistic avatar eye and head animation using a neurobiological model of visual attention. In *Proc. SPIE*.

JOHNSON, W. L., RICKEL, J. W., AND LESTER, J. C. 2000. Animated pedagogical agents: Face-to-face interaction in interactive learning environments. *International Journal of Artificial intelligence in education 11*, 47–78.

KING, D., ROWE, A., AND LEONARDS, U. 2011. I trust you; hence i like the things you look at: Gaze cueing and sender trustworthiness influence object evaluation. *Social Cognition 29*, 476–485.

KING, S. A. 2008. Animating speech in games. In *Motion in Games*. Springer.

KLEINKE, C. L. 1986. Gaze and eye contact: a research review. *Psychological bulletin 100*, 78.

KRÄMER, N., KOPP, S., BECKER-ASANO, C., AND SOMMER, N. 2013. Smile and the world will smile with you-the effects of a virtual agent's smile on users' evaluation and behavior. *Int. J. Hum.-Comput. Stud. 71*, 335–349.

KULMS, P., KRMER, N., GRATCH, J., AND KANG, S.-H. 2011. Its in their eyes: A study on female and male virtual humans gaze. In *Intelligent Virtual Agents*. Springer Berlin Heidelberg.

LANCE, B., AND MARSELLA, S. C. 2008. The relation between gaze behavior and the attribution of emotion: An empirical study. In *Intelligent Virtual Agents*, IVA '08.

LE, B. H., MA, X., AND DENG, Z. 2012. Live speech driven head-and-eye motion generators. *IEEE Transactions on Visualization and Computer Graphics 18*, 1902–1914.

LEE, S. P., BADLER, J. B., AND BADLER, N. I. 2002. Eyes alive. In *Computer graphics and interactive techniques*, SIGGRAPH '02.

LEICHTENSTERN, K., BEE, N., ANDRÉ, E., BERKMÜLLER, U., AND WAGNER, J. 2011. Physiological measurement of trust-related behavior in trust-neutral and trust-critical situations. In *Trust Management V*. Springer.

MA, X., AND DENG, Z. 2009. Natural eye motion synthesis by modeling gaze-head coupling. In *Virtual Reality Conference, 2009. VR 2009. IEEE*.

MARTÍNEZ-MIRANDA, J., JUNG, B., PAYR, S., AND PETTA, P. 2008. The intermediary agents brain: Supporting learning to collaborate at the inter-personal level (short paper). *Autonomous agents and multiagent systems*.

MASUKO, S., AND HOSHINO, J. 2007. Head-eye animation corresponding to a conversation for cg characters. In *Computer Graphics Forum*.

NAKANO, Y. I., REINSTEIN, G., STOCKY, T., AND CASSELL, J. 2003. Towards a model of face-to-face grounding. In *Association for Computational Linguistics*, ACL '03.

OOSTERHOF, N. N., AND TODOROV, A. 2008. The functional basis of face evaluation. *National Academy of Sciences 105*, 11087–11092.

PEARCE, W. B. 1974. Trust in interpersonal communication.

QUEIROZ, R. B., BARROS, L. M., AND MUSSE, S. R. 2008. Providing expressive gaze to virtual animated characters in interactive applications. *Comput. Entertain. 6*, 41:1–41:23.

REALE, M., LIU, P., AND YIN, L. 2011. Using eye gaze, head pose, and facial expression for personalized non-player character interaction. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*.

STEPTOE, W., STEED, A., ROVIRA, A., AND RAE, J. 2010. Lie tracking: social presence, truth and deception in avatar-mediated telecommunication. In *SIGCHI Conference on Human Factors in Computing Systems*, CHI '10.

STONE, M., AND OH, I. 2008. Modeling facial expression of uncertainty in conversational animation. In *Embodied communication in humans and machines*, ZiF'06.

THUE, D., BULITKO, V., SPETCH, M., AND ROMANUIK, T. 2011. A computational model of perceived agency in video games. In *Artificial Intelligence in Interactive Digital Entertainment (AIIDE)*.

TOMLINSON, R., AND BAHRA, P. 1986. Combined eye-head gaze shifts in the primate. ii. interactions between saccades and the vestibuloocular reflex. *Journal of neurophysiology 56*, 1558–1570.

TRUTOIU, L. C., CARTER, E. J., MATTHEWS, I., AND HODGINS, J. K. 2011. Modeling and animating eye blinks. *ACM Trans. Appl. Percept. 8*, 17:1–17:17.

VANHATUPA, J.-M. 2011. Guidelines for personalizing the player experience in computer role-playing games. In *Foundations of Digital Games*.

VRIJ, A., AND MANN, S. 2004. Detecting deception: The benefit of looking at a combination of behavioral, auditory and speech content related cues in a systematic manner. *Group Decision and Negotiation 13*, 61–79.

WANG, A., EMMI, M., AND FALOUTSOS, P. 2007. Assembling an expressive facial animation system. In *ACM SIGGRAPH symposium on Video games*.

WILLIS, M. L., PALERMO, R., AND BURKE, D. 2011. Social judgments are influenced by both facial expression and direction of eye gaze. *Social cognition 29*, 415–429.

WYLAND, C. L., AND FORGAS, J. P. 2010. Here's looking at you kid: Mood effects on processing eye gaze as a heuristic cue. *Social Cognition 28*, 133–144.

YEE, N. 2006. Motivations for play in online games. *CyberPsychology & Behavior 9*, 772–775.