

Looking might be deceiving:
How the looking behavior of a virtual robot
might influence its perceived trustworthiness

Remco Runge (s0715239)
Department of Artificial Intelligence
Radboud University Nijmegen
R.Runge@student.ru.nl

Supervisors:

R. van den Brule
Donders Center for Cognition
Radboud University Nijmegen

W.F.G. Haselager
Donders Center for Cognition
Radboud University Nijmegen

February 3, 2012

Abstract

Trust is an important factor in human to human interaction. It helps us cooperate with people we might not even know. In this study, we investigated the possibility of changing the perceived trustworthiness of a robot by changing its looking behavior. To make an estimate of someone's trustworthiness without prior interaction, people look at the appearance and movement of the interaction partner.

Participants played an iterative trust game on a computer with a movie of a virtual robot playing the interaction partner. The looking behavior of this robot was manipulated to either not avoid or avoid eye contact. The reciprocation rate was also manipulated (high/low). The investment rate of the participant was used as a dependent variable as well as an explicit trust rating given on a questionnaire after the experiment.

Analyses over the combination of all the rounds, showed a marginal significant effect of looking behavior on the investment rate. This effect was not present in analyses of only the first or last round of the iterative trust game. Looking behavior might therefore have an effect on the amount of perceived trustworthiness.

Interaction effects of looking behavior and reciprocation rate were present. A high reciprocating robot which avoided eye contact got both higher investments as well as a higher trust rating on the questionnaire, compared to a high reciprocating robot which did not avoid eye contact. This effect was not present for a low reciprocating robot. The right combination of looking behavior and performance could therefore lead to higher perceived trustworthiness.

Contents

1	Introduction	4
1.1	Trust	4
1.2	Trust game	4
1.3	Appearance and behavior	5
1.4	The experiment	6
1.5	Research questions	7
2	Method	8
2.1	Participants	8
2.2	Stimuli	8
2.3	Task	9
2.4	Questionnaire	9
3	Results	10
3.1	First round	10
3.2	Last round	10
3.3	All rounds	11
3.4	Questionnaire	12
4	Discussion	14
4.1	Trust without prior interaction	14
4.2	Trust with prior interaction	14
4.3	Trust in multiple interactions	15
4.4	Trust rating	16
5	Conclusion	18
A	Trial lists	21

1 Introduction

In the recent years, social robots have been gaining more and more importance in the field of robotics. Robots are starting to enter our homes and our daily life. The number and intensity of interactions between humans and robots will only keep increasing. Therefore, it is important to make these interactions as fluent as possible.

In human - human interaction, trust plays an important role in cooperation. Placing trust in other people is something we do every day, multiple times a day. We trust the mailman to deliver our mail without opening it, we trust the police to keep us safe and not use their guns to rob us, we trust our friends not to use the secrets we tell them against us. Without trust, society as we know it would not be possible. Trusting others is beneficial for humans, therefore we place trust in others, despite of the possible risks involved.

When social robots are becoming a more important part of our lives, trust could become an important factor in the interaction with these robots. We humans do not see robots as mere machines, we tend to anthropomorphize them even beyond their limited capabilities (Duffy, 2003). Like the trust we place in other humans, the amount of trust we place in a robot should correspond with the amount of trust that the robot ‘deserves’ without placing too much trust in the robot.

1.1 Trust

Trust is seen by most theories as the dyadic (interpersonal) relationship between two parties. In which one of the parties is the actor (the trustor), which places trust in a specific confederate (the trustee) to reach a certain goal (Simpson, 2007). The trustor relies on the trustee’s cooperation to get to the desired outcome or recourse.

A robot can become a trustee when a user places a certain amount of trust in the robot. This amount of trust should be carefully calibrated. Trusting a robot beyond its capabilities could have dangerous consequences. A robot that is made to lift heavy unbreakable boxes, might not be suited to lift a child. The child could get seriously injured when the robot would lift it in the same way as an unbreakable box.

On the other hand, the user could also place too little trust into a robot. While this may not be as dangerous, it could make the robot less effective or even useless. If you do not trust the box lifting robot enough, you might end up lifting the boxes yourself instead.

People try to estimate the amount of trust they might place in others mostly by looking at their appearance (Willis and Todorov, 2006) and by looking at their behavior (Axelrod, 1981). The appearance of a person is used to make an instant trust estimation even before interacting with the person. Within 100 ms people can decide whether a face is trustworthy or not (Willis and Todorov, 2006).

1.2 Trust game

To measure the amount of trust placed into a confederate (the trustee) by a participant (the trustor), the trust game was developed. In this game, a

participant is given a certain amount of money to invest into the confederate. The confederate is then given the choice to return none, some, or all of the invested money to the participant. The returned money will then be tripled or quadrupled before the participant receives it. If a participant does not trust the confederate at all, the participant will not invest any money since he might not get it back. When the participant does trust the confederate, he or she might invest some or all of the money he or she has. The amount of invested money by the participant can be related to the amount of trust that the participant puts into the confederate according to Berg et al. (1995).

When people interact in a trust game (Berg et al., 1995) with someone who has either a trustworthy or an untrustworthy face, they tend to take riskier choices when interacting with the trustworthy faced person (van 't Wout and Sanfey, 2008), compared to the choices they make in interaction with the untrustworthy faced person. Even while the reciprocation rate is the same (both persons are equally trustworthy on the sole basis of task performance).

While people initially use appearance to make an implicit trust estimation, over time, they incorporate the performance of the interaction partner to create a more accurate explicit trust estimation. After repeated interactions in an iterative trust game with both trustworthy and untrustworthy faced partners, people lower their investment into the trustworthy faced partner. They realize that the trustworthy faced partner has the same reciprocation rate as the untrustworthy faced partner and that therefore both partners should be trusted equally (Delgado et al., 2005).

Chang et al. (2010) not only manipulated the facial trustworthiness, but also the reciprocation rate of the interaction partner, while conducting an iterative trust game. They found an interaction effect such that a partner with an aligned facial trustworthiness and reciprocation rate was trusted more than a partner with an unaligned facial trustworthiness and reciprocation rate. .

1.3 Appearance and behavior

The research by Delgado et al. (2005) and Chang et al. (2010) show that while people at first make an implicit trust rating based on the appearance of the partner, they use repeated interaction to create a more accurate explicit trust rating based on the performance of the partner.

In an experiment with a mock airport customs checkpoint, it was shown that non-verbal behavior is an important cue for trustworthiness (Kraut and Poe, 1980). The mock travelers were more likely to be searched if they behaved in a certain way. Appearing nervous, avoiding eye contact and shifting their posture increased the likability of being searched.

According to (Marsi and Rooden, 2007), expressing certainty can be related to trust. They argued that a question and answering system which is able to display its confidence will gain more trust, than a system which lacks this ability. In their research they used a virtual talking head to answer questions. They adapted the eyebrow movements as well as the head movements of this virtual head, to express certainty or uncertainty. The guidelines they used to express certainty and uncertainty can be seen in Table 1.

Participants were asked how certain or uncertain they thought the system was about an answer. The results of this study suggested that humans can

	Eyebrows:	Head:
Certain:	- few movements - frown with new information	- few movements - nodding with new information
Uncertain:	- many (unnecessary) movements - raising eyebrows	- many (unnecessary) movements - sideward movement (shaking) with new information

Table 1: The certainty and uncertainty behavior as used by Marsi and Rooden (2007).

correctly recognize animated facial expressions and head movement as being certain, but that only head movements are a consistent cue for uncertainty.

The results found by Marsi and Rooden (2007) could be also applicable to a much simpler robotic face, since there is no big difference between the recognition of facial expressions by humans in humanlike faces and the recognition of facial expression in simple robotic faces (Schiano et al., 2000).

While changing the appearance of a robot could give an instant change in the trust estimation made by the user, it might not be very practical. Changing the movement of a robot could be an easier way to change the amount of trust that is placed into the robot. A robot could change its style of movement to give an indication of the amount in which it can be trusted. When the task performance of a robot is bad, while moving smoothly, people tend to place less trust into this robot compared to a robot with shaky movements and a bad task performance (van den Brule et al., In prep).

Changing the looking behavior, by means of its head movement, of a robot also influences its likability for most people (Wang et al., 2006). In the experiment by Wang et al. (2006), a robot tracked the head of the person whom it was interacting with. When the robot adapted an avoiding strategy for its head movement (it avoided eye contact), it was seen as more enjoyable opposed to strategies in which it tracked the face with or without smooth movement or when it adapted a strategy in which it just kept staring straight ahead.

Therefore we think that, by manipulating only the looking behavior of a robot through changes in its head movement, we can change the amount in which the robot is trusted.

1.4 The experiment

We adapted the iterative trust game that was used by Delgado et al. (2005) and Chang et al. (2010) to a computer experiment in which the role of confederate would be played by a virtual robot. Due to budget limitations, the money that could be invested by the participant was only virtual.

The robot’s looking behavior and reciprocation rate where manipulated. The robot either tried to avoid eye contact by looking away for about 80% of the time, or would not try to avoid eye contact and would only look away for about 20% of the time. These looking behaviors where combined with either a robot that did reciprocate half of the money it received in 80% of the cases, or a robot that did not reciprocate any money in 80% of the cases.

After the experiment, a questionnaire was used to measure explicit trust. Participants were asked to give each robot a trustworthiness rating on a scale

from 1 (not trustworthy at all) to 7 (very trustworthy).

1.5 Research questions

We developed the following research questions in line with previous research:

1. Do people invest more money into a robot that does not avoid eye contact, compared to the amount of money they invest in a robot that does avoid eye contact if there is no history of prior interaction between them?
2. Do people invest more money into a robot that does not avoid eye contact, compared to the amount of money they invest in a robot that does avoid eye contact if there is prior interaction between them?

These first two questions are the core of what we would like to research. We can answer the first question by looking at the effect of looking behavior (avoiding/not avoiding) on the investment rate in the first round. For the second question, we will investigate the effect of looking behavior on the investment rate in the last round. As an addition, we would also like to answer the following questions.

1. Do people give a robot that does not avoid eye contact a higher explicit trust rating, compared to the trust rating they give to a robot that does not look away?
2. Is there an interaction effect on the investment rate between the looking behavior of the robot and its reciprocation rate?
3. Is there an interaction effect on the explicit trust rating between the reciprocation rate and the looking behavior?
4. Do people invest more money into a robot which has a high reciprocation rate if there is prior interaction between them?
5. Do people give a robot that has a high reciprocation rate, a higher explicit trust rating, compared to the trust rating they give to a robot that does not look away?

By looking at the effect of looking behavior on the trust rating given on the questionnaire, we can answer question 1 of the additional questions. Question 2 of the additional questions can be answered by looking at the interaction effect between looking behavior and reciprocation on the investment rate in the first round, in the last round and in the combination of all of the rounds.

Question 5 and 6 are used as a control. It is expected that, in the same way as when the trust game would be played with only humans Delgado et al. (2005), participants will base their investment rate and their trust rating on the previous experience with the robots reciprocation rate. To answer these questions, we will look at the effect of reciprocation rate on the both the investment rate in the last round as well as the trust rating given on the questionnaire.

2 Method

2.1 Participants

118 participants (mean age = 19.35, SD = 2.925, 108 females, 10 males) took part in the experiment. All the participants were students of the Radboud University Nijmegen. As a reward for taking place in the experiment, they received course credits.

2.2 Stimuli

We used a virtual model of the robot TWENDY-one (Iwata and Sugano, 2009) to create stimulus movies with the help of the program Vizard. TWENDY-one is designed as an assistant for elderly people. Its appearance is robot like, but at the same time it also has some resemblance to humans and is therefore able to evoke an emphatic response.

Two different looking behaviors were animated in Vizard to create an animation of a robot that avoided eye contact, and one of a robot that did not avoid eye contact. The robot which did not avoid eye contact, looked straight ahead for about 80% of the time, while the robot which avoided eye contact only looked straight ahead for about 20% of the time. To prevent the robot from looking to unnatural, it made small gradual movements while looking straight ahead

During each trial, we also manipulated the reciprocation rate of the robot. Therefore each of the two movies (avoiding/not avoiding) would be shown twice for each reciprocation rate (high/low). These looking behavior/reciprocation rate combinations resulted in four ‘different’ robots. To make sure that each of these robots were recognizable and would be seen as individuals, we used four different skin colors (blue, green, yellow and purple). A frame of one of the movies can be seen in figure 1.

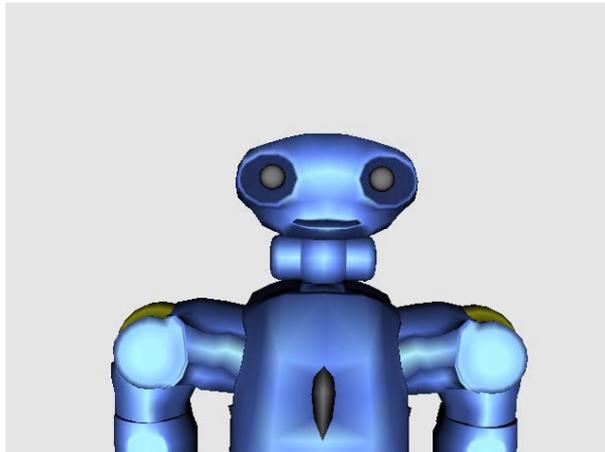


Figure 1: Frame from one of the stimuli movies.

2.3 Task

The participants played a trust game in which the part of the confederated was played by the computer. The trust game, as played in our experiment, went as follows. The participant would see one of the movies, with either a robot that avoided eye contact, or a robot that did not avoid eye contact. At the same time, the participant was given ten virtual euros and was asked to invest any amount of this money into the robot.

The participant was told that the invested amount would be quadrupled and then would be given to the robot. The participant was also told that the robot could choose to give back any amount of money. After investing money into the robot, the robot reciprocated half the money or no money at all depending on a pre-determined trial list (as can be found in appendix A.).

This process was repeated for all four of the reciprocation rate (low/high) and looking behavior (avoiding/not avoiding) combinations to form four trials, which in turn combined to one round of the trust game. The participants played a total of ten of these rounds. The investment rate of the participant was used as an implicit trust measure.

The trial list where created in such way, that over the ten rounds, a low reciprocating robot would have reciprocated 20% of the times, while a high reciprocating robot would have reciprocated 80% of the trials.

To be able to check for an effect of the sequence in which the movies where shown, two different sequences where used. The sequence which the participant would get for the entire experiment was counterbalanced.

To prevent effects of the color of the robots on the investment rate and on the trust rating we counterbalanced the colors of the robots between the participants.

2.4 Questionnaire

As an extra measurement of trust, we included a questionnaire at the end of our experiment. A picture of one of the robots looking straight ahead was shown above each of the questions. We asked the following questions (translated from dutch):

1. I found the robot trustworthy.
2. I found this robot unfriendly.
3. I found this robot easy to collaborate with.
4. I had the feeling that this robot made a connection with me.
5. This robot looked at me for ..% of the time.
6. This robot returned money in ..% of the cases.

Questions 1 to 4 where rated on a scale from "not at all" to "very much". Questions 5 and 6 where rated between 0 percent to 100 percent with intervals of 10 percent.

We were mostly interested in the first question. The other questions where used as control questions.

3 Results

Due to an error in the trial lists, the participants saw one of the stimuli movies twice during one of the rounds instead of seeing the four unique stimuli. For sequence 1 this happened in round 7, while in sequence 2 it happened in round 3. Therefore, we only used eight of the ten rounds in our analysis by excluding round 2 and round 7.

3.1 First round

The investment rate in the first round was examined with a 2x2x2 (sequence [1,2] x looking behavior [avoiding, non avoiding] x reciprocation rate [high, low]) repeated measures multivariate analysis of variance (ANOVA). Looking behavior and the reciprocation rate were within-subject factors, while sequence was a between subject factor.

The investment rate into a robot that does avoid eye contact is not significantly different from the investment rate into a robot that does not avoid eye contact. The effect of looking behavior was not significant: $F(1,85) = .829$. The participants did not invest more money into a robot that does avoid eye contact, compared to a robot that does not avoid eye contact if there is no history of prior interaction between them.

There was a significant difference in the investment rate between the two sequences in the first round: $F(1,85) = 7.900$, $p = .006$, $\eta_p^2 = 0.085$. The amount of invested money by sequence 1 was less ($M=3.163, SD=.186$) than the amount of invested money by sequence 2 ($M = 3.367, SD = .210$).

Since there is no prior history of interaction between the participants and the robots, we did not expect the reciprocation rate to have an effect in the first round. This expectation was confirmed by the results. There was no effect of reciprocation rate in the first round: $F(1,85) = 0.387$.

There was a significant interaction between the looking behavior and the reciprocation rate on the investment rate in the first round: $F(1,85) = 6.774$, $p = 0.011$, $\eta_p^2 = 0.005$. Participants invested more in a low reciprocating robot with avoiding looking behavior, compared to the amount they invested into a low reciprocating robot with not avoiding looking behavior $F(1,85) = 6.586$, $p = .012$, $\eta_p^2 = .072$. There was no significant difference between the investment rate into a high reciprocating robot with avoiding looking behavior, and a high reciprocating robot with not avoiding looking behavior $F(1,85) = 1.509$, $p > .05$.

Participants invested more money into a robot with avoiding looking behavior and a low reciprocation rate, compared to a robot with avoiding looking behavior and a high reciprocation rate: $F(1,85) = 5.309$, $p = .024$, $\eta_p^2 = .059$. The investment rate did not significantly vary between a robot with not avoiding looking behavior and a low reciprocation rate, and a robot with not avoiding looking behavior and a high reciprocation rate $F(1,85) = 1.737$ $p > .05$. These effects have been plotted in figure 2.

3.2 Last round

The investment rate in the last round was examined with a 2x2x2 (sequence [1,2] x looking behavior [avoiding, non avoiding] x reciprocation rate [high, low]) re-

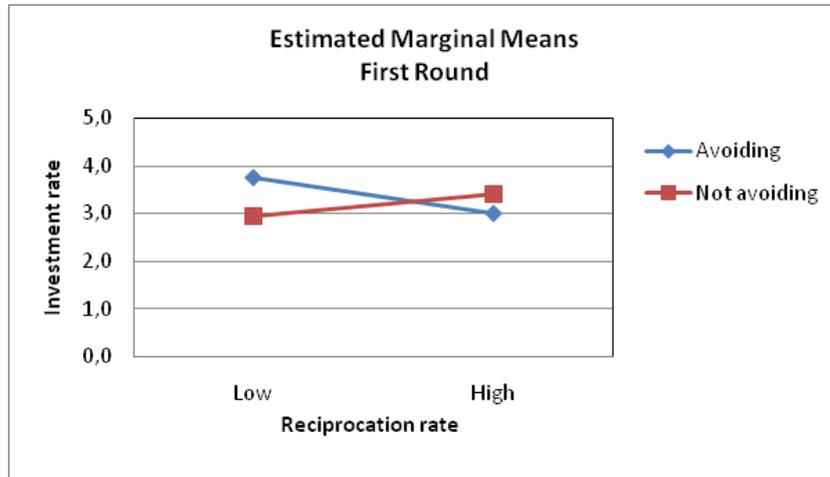


Figure 2: The estimated marginal means in the first round.

peated measures multivariate analysis of variance (ANOVA). Looking behavior and reciprocation rate were within-subject factors, while sequence was a between-subject factor.

The investment rate in a robot that does avoid eye contact was not significantly different from investment rate in a robot that does not avoid eye contact. The effect of looking behavior was not significant: $F(1,110) = .394$. The participants did not invest more money into a robot that did avoid eye contact, compared to a robot that did not avoid eye contact if there is prior interaction between them.

The investment rate into a robot with a high reciprocation rate was higher ($M = 3.554$, $SD = 2.229$), compared to the investment rate into a robot with a low reciprocation rate ($M = 2.888$, $SD = .234$): $F(1,110) = 8.008$, $p = .006$, $\eta_p^2 = .68$. The participants invested more money into a robot which had a high reciprocation rate if there was no prior interaction between them.

There was no interaction effect found between reciprocation rate and looking behavior on the investment rate in the last round $F(1,110) = .161$, $p > 0.05$, $\eta_p^2 = .001$.

3.3 All rounds

The investment rate in all of the rounds was examined with a $8 \times 2 \times 2 \times 2$ (round, sequence [1,2] x looking behavior [avoiding, non avoiding] x reciprocation rate [high, low]) repeated measures multivariate analysis of variance (ANOVA). Looking behavior and the reciprocation rate were within-subject factors, while sequence was a between-subject factor.

The trust rating given by the participants on the questionnaire was examined with a $8 \times 2 \times 2 \times 2$ (sequence [1,2] x looking behavior [avoiding, non avoiding] x reciprocation rate [high, low]) repeated measures multivariate analysis of variance (ANOVA). Looking behavior and the reciprocation rate were within-subject factors, while sequence was a between-subject factor.

The looking behavior of the robot had a marginal significant effect on the

amount of invested money: $F(1,74) = 3.915$, $p = 0.052$, $\eta_p^2 = 0.050$. Participants invested more into a robot that avoided eye contact ($M = 3.744$, $SD = 1.81$) than into a robot that did not avoid eye contact ($M = 3.528$, $SD = .166$).

The effect of reciprocation rate on the investment rate was significant: $F(1,74) = 64,717$, $p < .001$, $\eta_p^2 = .467$. Participants invested more money into a robot that had a high reciprocation rate ($M = 4.166$, $SD = .177$) compared to the amount of invested money into a robot that had a low reciprocation rate ($M = 3.105$, $SD = .179$).

There was a significant interaction between the looking behavior and the reciprocation rate in all of the rounds: $F(1,74) = 7,369$, $p = 0.008$, $\eta_p^2 = 0.091$. A robot with avoiding looking behavior and a high reciprocation rate, got a higher investment rate compared to a robot with avoiding looking behavior and a low reciprocation: $F(1,74) = 21.470$, $p < .001$, $\eta_p^2 = .225$. A robot which did not avoid eye contact with a high reciprocation rate, got a higher investment rate than a robot with not avoiding looking behavior and a low reciprocation : $F(1,74) = 58.612$, $p < .001$, $\eta_p^2 = .442$.

The investments in a robot with a low reciprocation rate and not avoiding looking behaviors, where not significantly different from the investments in a robot with a low reciprocation rate and avoiding looking behavior: $F(1,74) = .410$, $p = .524$, $\eta_p^2 = .006$. The participants did invest more into a robot with a high reciprocation rate and avoiding looking behavior, compared to a robot with a high reciprocation rate and not avoiding looking behavior: $F(1,74) = 8.177$, $p = .006$, $\eta_p^2 = .100$. These effects are plotted in figure 3.

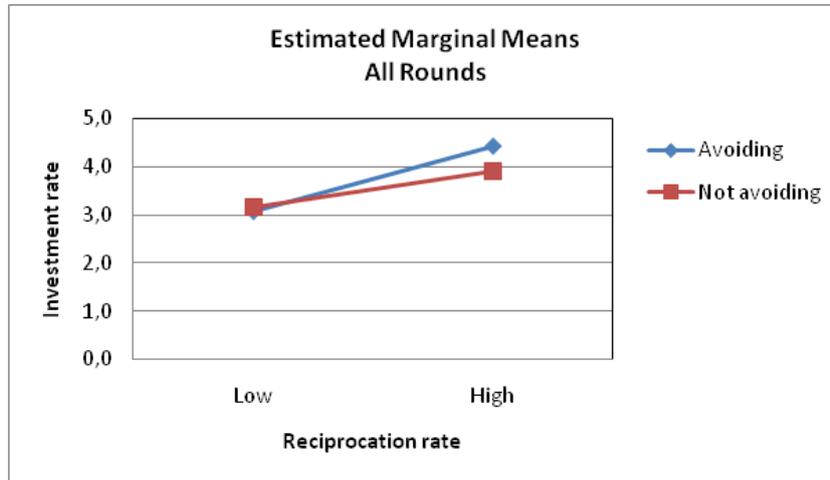


Figure 3: The estimated marginal means of all the rounds.

3.4 Questionnaire

In the questionnaire, the participants were asked how much they trusted each robot. There was, contrary to the marginal significant effect of the looking behavior on the investment rate, no significant effect on the amount of trust according to the questionnaire: $F(1,90) = 2.626$. The participants did not give

a robot that does avoid eye contact a higher explicit trust rating, compared to the trust rating they give to a robot that does not avoid eye contact.

The reciprocation rate of the robot had a significant effect on the trust rating given by the participants. A robot that had a high reciprocation rate received a higher trust rating ($M = 4.500$, $SD = .103$) compared to a robot which had a low reciprocation rate ($M = 3.000$, $SD = .113$): $F(1,90) = 73.563$, $p = .000$, $\eta_p^2 = .450$.

There was a significant interaction effect between the looking behavior of the robot and its reciprocation rate on the trust rating: $F(1,90) = 7.368$, $p = .008$, $\eta_p^2 = .076$. A robot which avoided eye contact with a high reciprocation rate, got a higher trust ratings than a robot which avoided eye contact with a low reciprocation rate $F(1,90) = 19.688$, $p < .001$, $\eta_p^2 = .179$. The same was true for a robot that did not avoid eye contact. A robot which did not avoid eye contact with a high reciprocation rate, got a higher trust ratings than a robot which also did not avoid eye contact with a low reciprocation rate $F(1,90) = 72.728$, $p < .001$, $\eta_p^2 = .447$.

There was no significant difference in the investments into a low reciprocating robot with an avoiding looking behavior and the investments into a low reciprocating with not avoiding looking behavior: $F(1,90) = .729$, $p = .395$, $\eta_p^2 = .008$. The participants did invest more into a high reciprocating robot with avoiding looking behavior, compared to a high reciprocating robot with not avoiding looking behavior: $F(1,90) = 8,517$, $p = .004$, $\eta_p^2 = .086$. These effects are displayed in figure 4.

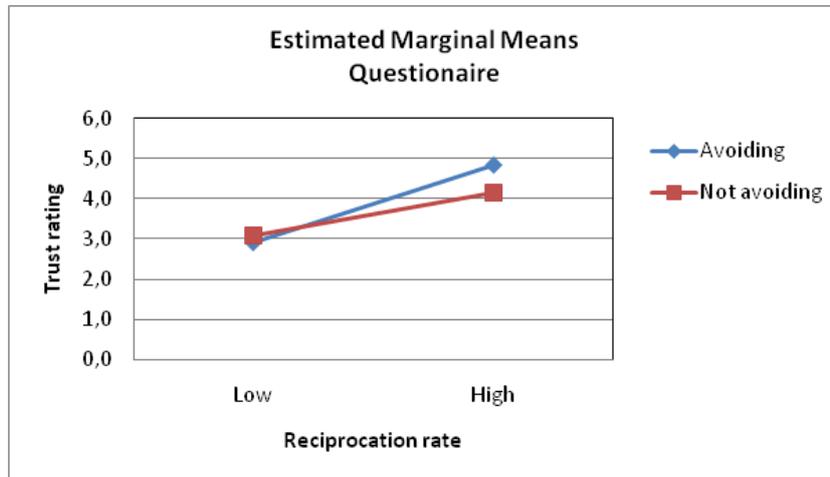


Figure 4: The estimated marginal means of the questionnaire.

4 Discussion

In this bachelor thesis, we investigated the possibility to manipulate the perceived trustworthiness of a robot by changing its looking behavior. Previous research has shown that initial impressions plays an important role in judging someone’s trustworthiness (Delgado et al., 2005; van ’t Wout and Sanfey, 2008). With the help of an iterative trust game (Berg et al., 1995; Chang et al., 2010), we studied the amount of trust placed into a virtual robot.

Our first main research questions was, if a robot with avoiding looking behavior would get a lower investment rate, compared to a robot that has a not avoiding looking behavior in the first round. Our second main research question was, if people invest more money into a robot that does not avoid eye contact, compared to amount of money they invest in a robot that does avoid eye contact if there is prior interaction between them.

4.1 Trust without prior interaction

Unfortunately, we did not find a significant main effect of looking behavior on the investment rate in the first round. The invested amount into the robot did not vary between the two looking behaviors.

One reason for this might be that the participants did not have any previous experience with the trust game before the start of the first round. Therefore, they might have been concentrating too much on the control of the experiment as well as the text, instead of looking at the movies.

It is also possible that the sequence in which the looking behaviors combined with the reciprocation rates were presented to the participants, has had an influence on the investment rate of the participant. This is underlined by the significant main effect of sequence that was found. If the first movie that was shown did not reciprocate, the participants might adapt their strategy in fear of also losing their money to the next robot. For example, the participants with sequence 1 got a high reciprocating robot in the first trial, while the participants in sequence 2 got a low reciprocating robot in the first round. The investments by the participants with sequence 1 on the second trial was on average higher ($M = 3.00$, $SD = 2.200$) compared to the investments by the participants with sequence 2 on the second trial ($M = 2.67$, $SD 2.220$).

We did find an interaction effect between the looking behavior and reciprocation rate in the first round. This effect was unexpected since there was no prior interaction between the participant and the robot. Therefore the participant could not adept its investments based on the reciprocation of the robot. The interaction effect might have been caused by the sequence in which the stimuli where presented to the participant. If the first robot would have had a high reciprocation rate, the participant might expect the next robot to also have a high reciprocation rate, as explained in the previous paragraph.

4.2 Trust with prior interaction

Do people invest more money into a robot that does not avoid eye contact, compared to a robot that does avoid eye contact if there is prior interaction between them? We did not find a significant main effect of looking behavior in the last round. The participants did not invest more money into a robot that

did not avoid eye contact compared to one that did avoid eye contact. This finding is what we expected based on previous research (Delgado et al., 2005).

After multiple interactions, people start to base their trust onto the reciprocation rate of the robot, instead of its appearance. This is supported by the significant main effect of reciprocation rate on the investment rate in the last round. The participants had a higher investment rate into a robot with a high reciprocation rate, compared to the investment rate into a robot with a low reciprocation rate. The participants learned after several iterations, which robot would yield a high return, and which robot would not.

We did not find a significant interaction effect of looking behavior and reciprocation rate in the last round. Our manipulation might not have been strong enough to yield a significant interaction effect.

4.3 Trust in multiple interactions

While we did not find a significant effect of looking behavior on the investment rate in the first and in the last round, we did find a marginal significant effect of the looking behavior if we look at all the rounds. The participants invested more money into a robot that avoided eye contact, compared to a robot that did not avoid eye contact. This effect is the other way round from what we expected based on the experiment by Kraut and Poe (1980). They found that people who avoid eye contact are seen as being less trustworthy. Though the research of Wang et al. (2006) showed that a robot which avoided eye contact was seen as being more likable compared to robots whom did not avoid eye contact.

The not avoiding eye contact robot might have made the participants feel uneasy by staring too much. This could have decreased the perceived trustworthiness of the robot. Another explanation could be that the not avoiding eye contact robot did not move naturally enough. It moved less than the avoiding eye contact robot and therefore might have been too unnatural to interact with. Participants might trust a natural, or more human like, moving robot more than a robot which is too unnatural in its movement. The avoiding looking behavior might also have caused the participants to think that the robot does not trust them, and therefore invest more money into the avoiding robot in an effort to gain its trust.

Like the interaction effect of reciprocation rate and looking behavior in subsection 4.1, the effect of looking behavior might also have been caused by the sequence in which the behaviors were presented to the participants.

As in the analyses of last round, we found the expected effect of reciprocation rate on the investment rate if we looked at all the rounds. Participants invested more money into a robot with a high reciprocation rate compared to a robot with a low reciprocation rate in all of the rounds. As described in the previous subsection (4.2) participants probably learned in which robot to invest to get the highest return.

There was a significant interaction between the looking behavior and the reciprocation rate in all of the rounds on the investment rate. A high reciprocating robot with avoiding looking behavior, got a higher investment rate than a robot with a high reciprocation rate and a not avoiding looking behavior. There was no such effect present in case of a low reciprocating robot. This effect might also have been caused by the sequence in which the stimuli were presented. As

the effect of looking behavior, the lower investments into the robot that does not avoid eye contact might have been caused by it staring too much, or moving to unnatural. The higher investment rate for the avoiding robot might have been caused by the overcompensation of the participants in an effort to gain the robots trust.

Participants invested more money into a high reciprocating robot with either avoiding or not avoiding looking behavior than into a low reciprocating robot with either avoiding or not avoiding looking behavior. This was probably caused in the same way as the main effect of reciprocation. Participants learned to invest in the robots that would give the highest returns.

4.4 Trust rating

With the help of a questionnaire, we also tried to measure the explicit trust the participants placed into the robots. There was no significant effect found of the looking behavior on the explicit trust rating on the questionnaire. This is what we expected based on the research of (Delgado et al., 2005). The questionnaire was at the end of the test, and therefore the participants might have based their explicit trust rating on the reciprocation rate instead of the looking behavior due to the repeated interaction.

The reciprocation rate did have a significant effect on the explicit trust measure in the questionnaire. The participants placed more trust in a robot with a high reciprocation rate, compared to amount of trust they reported placing into a robot with a low reciprocation rate. Participants probably learned which robots to invest in to get the highest return, which caused this effect as described in the previous subsections.

The looking behavior and the reciprocation rate had an interaction effect on the trust rating. A high reciprocating robot with avoiding looking behavior, got a higher investment rate compared to a high reciprocating robot with not avoiding looking behavior. For a low reciprocation with either avoiding or not avoiding looking behavior, the investment rate did not vary. As explained in the previous subsection this effect might have been caused by the participants feeling uneasy due to the staring of the not avoiding robot. The movements of the avoiding robot might also have been too unnatural to gain the participants trust. The avoiding robot might have received higher investments due to the overcompensation of the participants in an effort to win the robots trust. The effect might have also have been caused by the sequence in which the trials were presented to the participants.

A robot which avoided eye contact with a high reciprocation rate, got a higher trust rating than the trust rating given to an avoiding eye contact robot with a low reciprocation rate. For a robot which did not avoid eye contact, the same was true. When this robot had a high reciprocation rate, it led to a higher trust rating, than the trust rating given to the robot with the same looking behavior and a low reciprocation rate. As in the previous sections, this effect can be explained by the participants learning which robots have high reciprocation rates and are therefore receive higher trust ratings..

This interaction effect looks a lot like the interaction effect we found on the investment rate by the looking behavior and reciprocation rate. There might be a connection between these two interaction effects. Due to the scope of our

research, we did not analyze this effect. For future research, this could be done with a covariate analyses.

5 Conclusion

We showed that people will adapt their investment behavior based on the reciprocation rate of a robot after playing multiple iterations of a trust game. While we were unable to find an effect of looking behavior on the investment rate in the first and last round, we did find a marginal significant effect of the looking behavior of the robot on the investment rate by the participants. This shows that the perceived trustworthiness of a virtual robot might be influenced by the looking behavior of that robot.

We did also find interaction effects of the looking behavior and the reciprocation rate on both the investment rate, as well as the trust rating. Therefore, the right combination of looking behavior and reciprocation rate of a virtual robot, could lead to a higher perceived trustworthiness. Although this effect might only be present if the robot has a high performance.

Further research is needed to come to a more definitive conclusion about the effects of looking behavior on the amount of perceived trustworthiness of a virtual robot. Different types of looking behavior, as well as different types of body language could be investigated to aid the development of social robots. Looking behavior in human to human interaction might be studied to create a more human like looking behavior for the (virtual) robot in an effort to increase the perceived trustworthiness of a robot.

In the end, without any previous interaction, we should be able perceive which robot to trust, and which we should not.

References

- R Axelrod. The evolution of cooperation. *Science*, 211:1390–1396, 1981. URL <http://www.sciencemag.org/content/211/4489/1390.short>.
- J. Berg, J. Dickhaut, and K. McCabe. Trust, reciprocity, and social history. *Games and economic behavior*, 10(1):122–142, 1995. URL <http://www.sciencedirect.com/science/article/pii/S0899825685710275>.
- Luke J Chang, Bradley B Doll, Mascha van 't Wout, Michael J Frank, and Alan G Sanfey. Seeing is believing: trustworthiness as a dynamic belief. *Cognitive psychology*, 61(2):87–105, September 2010. ISSN 1095-5623. doi: 10.1016/j.cogpsych.2010.03.001. URL <http://www.ncbi.nlm.nih.gov/pubmed/20553763>.
- M R Delgado, R H Frank, and E a Phelps. Perceptions of moral character modulate the neural systems of reward during the trust game. *Nature neuroscience*, 8(11):1611–8, November 2005. ISSN 1097-6256. doi: 10.1038/nn1575. URL <http://www.ncbi.nlm.nih.gov/pubmed/16222226>.
- Brian R. Duffy. Anthropomorphism and the social robot. *Robotics and Autonomous Systems*, 42(3-4):177 – 190, 2003.
- H. Iwata and S. Sugano. Design of human symbiotic robot twenty-one. In *Robotics and Automation, 2009. ICRA '09. IEEE International Conference on*, pages 580 –586, may 2009.
- Robert E Kraut and Donald Poe. Behavioral Roots of Person Perception : The Deception Judgments of Customs Inspectors and Laymen. *Journal of Personality and Social Psychology*, 39(5):784–798, 1980.
- Erwin Marsi and Ferdi Van Rooden. Expressing uncertainty with a talking head in a multimodal question-answering system, 2007.
- Diane J. Schiano, Sheryl M. Ehrlich, Krisnawan Rahardja, and Kyle Sheridan. Face to interface: facial affect in (hu)man and machine. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, CHI '00, pages 193–200, New York, NY, USA, 2000. ACM.
- Jeffrey a. Simpson. Psychological Foundations of Trust. *Current Directions in Psychological Science*, 16(5):264–268, October 2007. ISSN 0963-7214. doi: 10.1111/j.1467-8721.2007.00517.x. URL <http://cdp.sagepub.com/lookup/doi/10.1111/j.1467-8721.2007.00517.x>.
- R. van den Brule, R. Dotch, G. Bijlstra, D.H.J. Wigboldus, and W.F.G. Haselager. A shaky foundation for trust: Effects of task performance and movement style on trust and behavior in social human-robot interaction, In prep.
- M van 't Wout and a G Sanfey. Friend or foe: the effect of implicit trustworthiness judgments in social decision-making. *Cognition*, 108(3):796–803, September 2008. ISSN 0010-0277. doi: 10.1016/j.cognition.2008.07.002. URL <http://www.ncbi.nlm.nih.gov/pubmed/18721917>.

Emily Wang, Constantine Lignos, Ashish Vatsal, and Brian Scassellati. Effects of head movement on perceptions of humanoid robot behavior. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction, HRI '06*, pages 180–185, New York, NY, USA, 2006. ACM.

Janine Willis and Alexander Todorov. Making Up Your Mind After a 100-Ms Exposure to a Face. *Psychological Science*, 17(7):592–598, 2006.

A Trial lists

On the next pages, you can find the two trial lists as used in the experiment.

Round	Trial	Looking behavior	Reciprocation rate	Reciprocate
1	1	Not avoiding	High	True
1	2	Not avoiding	Low	False
1	3	Avoiding	High	True
1	4	Avoiding	Low	False
2	5	Avoiding	High	True
2	6	Not avoiding	Low	False
2	7	Avoiding	High	True
2	8	Avoiding	Low	False
3	9	Not avoiding	High	False
3	10	Not avoiding	Low	True
3	11	Avoiding	High	True
3	12	Avoiding	Low	False
4	13	Not avoiding	High	True
4	14	Not avoiding	Low	False
4	15	Avoiding	High	False
4	16	Avoiding	Low	True
5	17	Not avoiding	High	True
5	18	Not avoiding	Low	False
5	19	Avoiding	High	True
5	20	Avoiding	Low	False
6	21	Not avoiding	High	True
6	22	Not avoiding	Low	False
6	23	Avoiding	High	True
6	24	Avoiding	Low	False
7	25	Not avoiding	High	True
7	26	Not avoiding	Low	False
7	27	Avoiding	High	True
7	28	Avoiding	Low	False
8	29	Not avoiding	High	False
8	30	Not avoiding	Low	True
8	31	Avoiding	High	True
8	32	Avoiding	Low	False
9	33	Not avoiding	High	True
9	34	Not avoiding	Low	False
9	35	Avoiding	High	False
9	36	Avoiding	Low	True
10	37	Not avoiding	High	True
10	38	Not avoiding	Low	False
10	39	Avoiding	High	True
10	40	Avoiding	Low	False

Table 2: Trial list for sequence 1

Round	Trial	Looking behavior	Reciprocation rate	Reciprocate
1	1	Avoiding	Low	False
1	2	Not avoiding	Low	False
1	3	Avoiding	High	True
1	4	Not avoiding	High	True
2	5	Avoiding	Low	False
2	6	Not avoiding	Low	False
2	7	Avoiding	High	True
2	8	Not avoiding	High	True
3	9	Avoiding	Low	False
3	10	Not avoiding	Low	True
3	11	Avoiding	High	True
3	12	Not avoiding	High	False
4	13	Avoiding	Low	True
4	14	Not avoiding	Low	False
4	15	Avoiding	High	False
4	16	Not avoiding	High	True
5	17	Avoiding	Low	False
5	18	Not avoiding	Low	False
5	19	Avoiding	High	True
5	20	Not avoiding	High	True
6	21	Avoiding	Low	True
6	22	Not avoiding	Low	False
6	23	Avoiding	High	True
6	24	Not avoiding	High	False
7	25	Avoiding	Low	False
7	26	Not avoiding	Low	False
7	27	Avoiding	High	True
7	28	Avoiding	High	True
8	29	Avoiding	Low	False
8	30	Not avoiding	Low	True
8	31	Avoiding	High	True
8	32	Not avoiding	High	False
9	33	Avoiding	Low	False
9	34	Not avoiding	Low	False
9	35	Avoiding	High	False
9	36	Not avoiding	High	True
10	37	Avoiding	Low	False
10	38	Not avoiding	Low	False
10	39	Avoiding	High	True
10	40	Not avoiding	High	True

Table 3: Trial list for sequence 2