Do people predict more rationality for an agent who is in a hurry?

Bachelor Thesis

Loes Immens
Student number: 0742562

External supervisor: Drs. Markus Paulus
Internal supervisor: Dr. Iris van Rooij
External supervisor: Dr. Sabine Hunnius

17 June 2011

Abstract

In many situations, humans are able to predict other agents' behavior. It would be useful if a robot would also be able to do this. Cognitive models can be implemented in a robot. For robots and cognitive models, humans are used as inspiration. In many cognitive models about action prediction, *rationality,* executing the action with the best possible outcome under the constraints of the situation, is assumed (Baker, Tenenbaum, & Saxe, 2007; Blokpoel, Kwisthout, van der Weide, & van Rooij, 2010). However, do humans really expect other agents to act rationally? In other studies (Gergely, Nádasdy, Csibra and Bíró, 1995; Csibra, Gergely, Bíró, Kóos, and Brockbank, 1999; Paulus, Hunnius, van Wijngaarden, Vrins, van Rooij, and Bekkering, in press) rationality was entangled with frequency effects. In the present study, we excluded frequency information. We expected adults to predict other agents' actions to be rational in the sense of efficient and we expected that this effect was stronger in a race than an everyday context. However, our results suggest that people are not efficient and that the race context does not influence this effect.

Imagine you are sitting in a restaurant. Sometimes, if it is not too busy and it is a good restaurant, the waiters come to pick up your empty plates right when you finish your course, and return with the next course almost seamlessly after that. The staff of such a restaurant is well attuned to each other and to their guests. They predict when the guests will finish their course and they adapt their actions to that.

It seems that to a certain extent, humans are able to predict other agents' actions or needs. However, it would be useful if this might also be done by a robot. For example, a person with a disability would benefit from a robot that could help them with daily activities. This would make them less dependent on other people.

A subfield of artificial intelligence that studies how robots can communicate with or assist humans is called social robotics. *Social robots* are robots that interact with each other and with their environment. Social robots are autonomous (Heerink, 2006) and embodied (Dautenhahn & Billard, 1999), which means that they are able to operate on their own and that they have a physical body with which they interact with their environment. Social robots are capable of social learning and imitation, are able to recognize their interaction partner (Dautenhahn & Billard, 1999; Fong, 2003) and to communicate with gestures and natural language (Fong, 2003). Social robots can have different functions. Some social robots function as assistants or servants to humans (Fong, 2003). Example of these are robot vacuum cleaners (Forlizzi & DiSalvo, 2006; Prassler, Ritter, Schaeffer, & Fiorini, 2000). Others are used as companions or pets (Fong, 2003). Examples are the iCat (van Breemen, Yan, & Meerbeek, 2005), AIBO (Fujita, 2001) and Paro (Wada, Shibata, Shaito, & Tanie, 2003).

In order to communicate with their environment, social robots need a cognitive representation of that environment. *Cognitive computational modeling* is a subfield in artificial intelligence in which computational models are designed for cognitive functions (Baker, Saxe, & Tenenbaum, 2009; Baker, Tenenbaum, & Saxe, 2006). These models can be implemented in a robot. Both in the fields of social robots and cognitive computational modeling, humans are a great source of inspiration. If humans are good at some task, an obvious way to create a robot that is also good at this task, is to study the way humans deal with the task and try to simulate it. One thing that is often assumed in cognitive models is rationality (Baker, Tenenbaum, & Saxe, 2007; Blokpoel, Kwisthout, van der Weide, & van Rooij, 2010). The most *rational* action for an agent is the action that maximizes the reward for the agent, while minimizing the costs, under the constraints of the situation. For example, efficiency is rational if you want to hurry. Because the assumption that humans predict other agents to act rationally is widespread in cognitive computational modeling (Baker et al., 2007; Blokpoel et al., 2010), this raises the question if humans really expect other agents to behave rationally, or not.

Gergely, Nádasdy, Csibra and Bíró (1995) did a habituation study to investigate whether 9-month-olds and 12-month-olds predict that others' actions are rational, which was repeated by Csibra, Gergely, Bíró, Kóos, and Brockbank (1999) and many others. In these studies, the infants were shown a video in which a small circle moved towards a big circle (see Figure 1). In the experimental condition, the circles were separated by an obstacle. The small circle had to jump over the obstacle to reach the big circle. In the control condition, the obstacle was not between the circles. However, the small circle made the same jumping movement. In the test phase, there was no obstacle. In the 'old action', the small circle jumped, like before. In the 'new action', the small circle took the direct route to the big circle. Gergely et al. (1995) and Csibra et al. (1999) saw the 'new action' as more rational, because the goal was reached more efficiently. They found that the infants in the experimental condition looked longer at the 'old action', and interpreted this as surprise that the small circle did not perform the most efficient action. These experiments were used as evidence for the *principle of rational action* (Gergely & Csibra, 2003), according to which people predict that others try to reach a goal by using the most efficient actions under situational constraints.

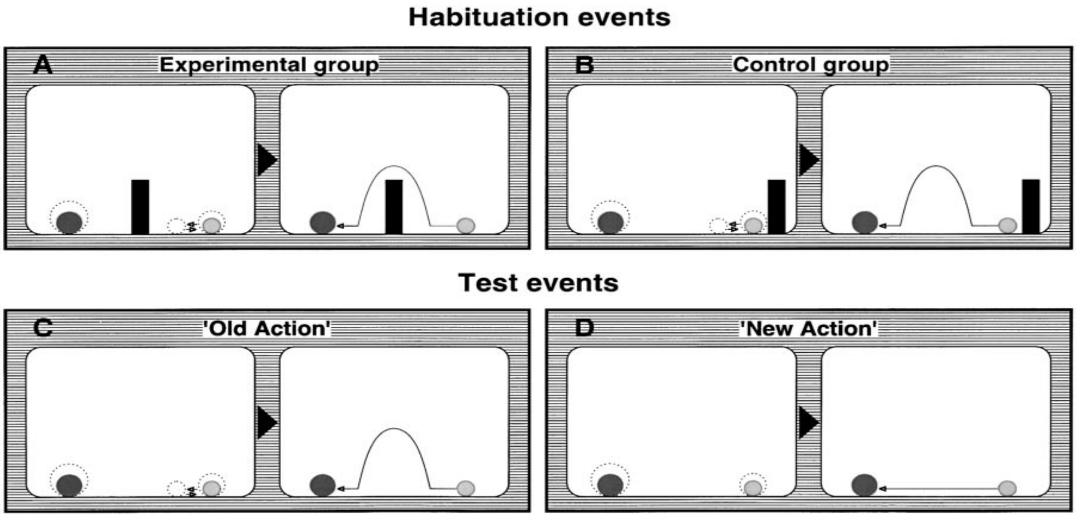

Figure 1: The experiments of Gergely et al. (1995) and Csibra et al. (1999)

However, Paulus, Hunnius, van Wijngaarden, Vrins, van Rooij, and Bekkering (in press) proposed an alternative explanation for these results. Because of statistical learning, the ‘new action’ is more familiar to infants. In watching and experiencing what happens around them in their daily lives, infants learn that most agents do not jump suddenly, when moving over a flat surface,. Because of their experience, they might be more surprised to see the ball jump through the air in the ‘old action’ video than to see the ball moving straight over the floor in the ‘new action’ video. In other words, a *frequency effect* might explain their looking longer at the ‘old action’.

To disentangle frequency and rationality in predicting other agents’ actions, Paulus et al. (in press) designed a study in which they manipulated the frequency effect. 9-month-olds and adults were shown a video while eye tracking information was gathered. A cow and a sheep were on two sides of a screen, with a long path and an interrupted short path between them (see Figure 2a). The cow had to walk toward the sheep. Over the junction, at which the cow could go either way, an occluder was placed. This occluder started out transparent, to show there was a road under it. It gradually turned opaque. In the habituation phase, the cow always took the long path. In the test phase, there was no gap in the short path. In the 'old action' (see Figure 2b), the cow took the long path. In the 'new action' (see Figure 2c), the cow took the short path. While the cow was under the occluder, eye movements of the participants were measured, to see whether the participant anticipated to either of the paths.

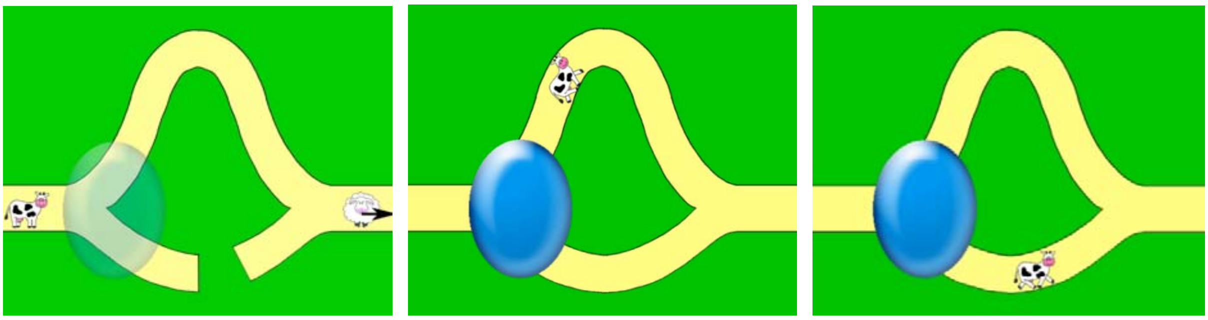

Figure 2: Some frames of the stimuli of Paulus et al. (in press)

(a) habituation phase **(b)** 'old action' **(c)** 'new action'

Paulus et al. (in press) found that in the first test trial, participants anticipated to the long path, which suggests that infants and adults use frequency information. However, after adults had seen the cow take the short path, they anticipated to the short path. A possible explanation is that adults use frequency information as a default, but switch to rationality-based predictions if expectations based on frequency are not met. Another explanation is that, in the first trial, adults overlooked the possibility of the cow taking the short path. With these results, Paulus et al. (in press) did not disentangle rationality and frequency completely. Also, different results were found for infants and adults. The question remained how adults would predict other agents' actions if there would be no frequency information available. In such a situation, would they base their expectations on rationality?

In the present study, we tried to exclude frequency learning by making it a one trial study, without a habituation phase. Eye movement information was gathered while adult participants were looking at the stimuli. A video showed a car driving on a road. At one point, there was a junction at which the car could take a short road or a long road. The short road would be the most rational road in the sense of efficiency. An occluder hid the car while it was on the junction. This occluder was called the *critical* occluder. However, whenever in the rest of the paper the word occluder is used, without specifying which occluder is meant, it is always about the critical occluder. Before the junction in the road appeared, there were several other occluders on the road. These were placed at straight parts and bends of the road to teach participants that the car reappears always stayed on the road and always reappeared from under the occluder. For these occluders, the car reappeared at the time when you would expect to reappear. The car spent more time under the critical occluder than expected to evoke participants' anticipation to either of the roads when the car did not appear at the expected time. We were also interested in creating a situation where the car would be expected to take the most efficient route. For this reason, there were two conditions. In the *race condition*, the car was in a race, implying that it would want to go as fast as possible. In the *control condition*, there was no race. The road and the speed of the car were the same in both conditions.

The questions we wanted to answer were: *Do people predict other agents' actions assuming that they behave rationally? Is this dependent on the context?* It is important to answer these questions, because rationality is assumed in many cognitive models about action prediction (Baker et al., 2007; Blokpoel et al., 2010). The answers to these questions might suggest whether this is justified.

We expected that participants would anticipate first to and look longer at the short road, which is the most efficient road, because the results of Gergely et al. (1995) and Csibra et al. (1999) suggested that rationality was assumed. In this study, we used efficiency as a measure of rationality. Also, the results of Paulus et al. (in press) suggested that participants predicted other agents' behavior based on the assumption of rationality, after expectations based on frequency information were not met. Because frequency learning was excluded in the present study, we expected that participants would predict the car's actions based on an assumption of rationality.

We expected that this effect would be stronger in the race condition than the control condition. The race context suggested that the car was in a hurry. In such a situation, efficiency would be more rational, so the short road would be more attractive, because it takes less time than the longer road.

Method

Participants

The participants were 27 adults (19 females and 8 males), recruited at the Radboud University in Nijmegen. Their ages ranged from 19 to 63 (M = 24.93). An additional 2 participants were excluded from the study. For these participants, calibration did not go well enough, causing shaky gaze movements. As a reward for participating in this study, participants received a chocolate bar.

Material

Tobii. The eye movements of the participants were tracked with a corneal reflection eye tracker (Tobii 1750, Tobii Technology, Sweden), which was integrated in a 17" TFT flat-screen monitor. On this monitor, the stimuli were shown.

Stimuli. The stimuli consisted of four movies, which were created in Adobe Image Ready 7.0 and Window Live Movie Maker 2011. Each movie had a length of 1 minute and 44 seconds and a resolution of 1280 x 1024 pixels. For both two conditions, there was a movie where the long road was on the right (*long right*) and one where the long road was on the left (*long left*).

Each movie started with an overview of the whole road from the start to the end (see Figure 3a). Then, the frame of the start was shown (see Figure 3b). In this frame, the car started moving. After the start, the road made several bends. In these bends and straight parts of the road (see for example Figure 3c and 3d), there were 6 occluders before the critical part of the road (see Figure 3e and 3f). The occluders hid the car while the car was under them. The car stayed under the occluders no longer than could be expected. This time was different for bends and straight parts in the road (400-1000ms). In the critical part of the road, the car stayed under the occluder longer than could be expected (3000ms). In this paper, the time when the car is under the occluder is defined as the period when the car is not visible at all. When the car reappeared, it took the long road. After that, the road ended and the car stopped moving.

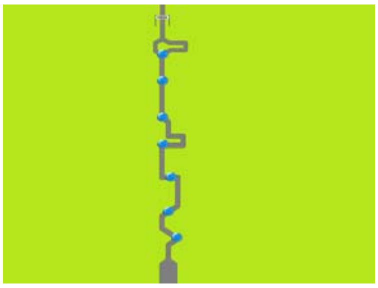

(a) Overview of the road

(b) Start

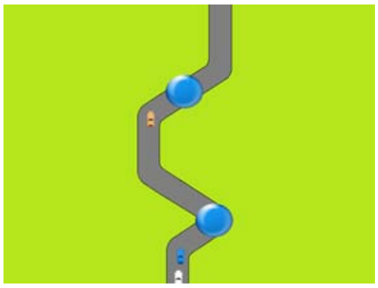

(c) The car is tailed by other cars

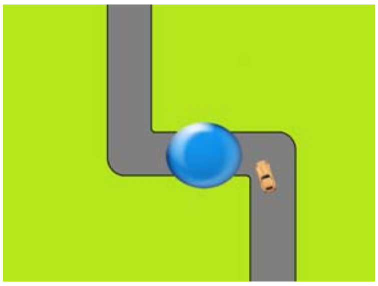

(d) The car is going towards an occluder

(e) The critical point, where the road splits into a short road and long road

(f) Finish

Figure 3: Some frames of the long right stimuli in the racing condition.

Race condition versus control condition.

The stimuli in the race and control condition resembled each other with the following differences. A finish sign was visible in the overview frame. In the start frame, there were 8 differently colored cars and racing lights in the racing condition, and only the yellow car and no lights in the control condition. First, the racing lights were out, then each light went on in turn from left to right, until finally all the lights went out again and the cars started moving. The yellow car went faster than the other cars. From then until the finish, there were no differences between the race condition and the control condition, except that in the race condition, sometimes you see a blue or a white car tailing the yellow car at the bottom of the screen. The last frame included a finish banner in the race condition, as opposed to in the control condition.

Questionnaire. In a short questionnaire some demographic questions were asked. Additionally, the questionnaire included an open ended question in which the participants were asked what they thought the experiment was about. This was to see whether they understood what we wanted to research. This knowledge might have influenced the results, because we used an implicit measurement.

Procedure

The experiment was run using the program Clearview (Tobii Technology, Sweden). The participants were instructed to watch a movie, while their eye movements were registered with a corneal reflection eye tracker (Tobii 1750, Tobii Technology, Sweden). They were seated at approximately 60 cm from the eye tracking device. They were instructed to sit as still as possible. The screen of the eye tracking device showed whether the eyes of the participant were visible. If not, the screen or the participant’s position was adjusted so that the device could detect the participant’s eyes. If the eyes were visible, the experimenter turned off the lights. The participants were instructed to follow 9 calibration points with their eyes. If any calibration points were not calibrated, these were repeated. If, after three tries, there were still calibration points that were not calibrated, calibration was aborted. After the calibration phase, the movie was started. Participants were asked to look at the general screen area, but received no more specific instruction. After the movie, the light was turned on again. Participants answered a short questionnaire.

Data analysis

Areas of interest (AOI) of equal shape and size and equal distance from the middle of the road were specified for both roads (see Figure 4). If a participant looked at the AOI of a road, this was scored as the participant looking at that road. In the rest of this paper, if a participant is said to have looked at a certain road, this means that the participants looked at the AOI of that road.

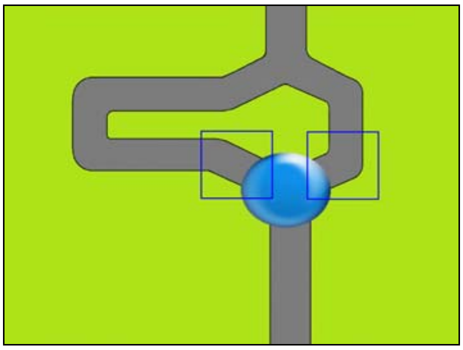

Figure 4: The areas of interest (AOIs)

For the analyses, we did not use the raw data. In Clearview, we specified the period when the car was under the occluder. Also, we specified that if a participant looked at a certain AOI for at least 100 ms, that was considered a gaze. In this way, we received a file for every participant, showing for every gaze to an AOI how many milliseconds that participant had looked at that AOI.

The dependent variables in this study were called anticipation and looking time. While the car was under the occluder, the road to which the participant looks first was called the *anticipation* of a participant. Anticipation could have three values: short road, long road or no anticipation. The short and the long roads were represented by -1 and 1, while no anticipation was represented by 0.

Looking time was a variable computed from the *looking time for the short road* and the *looking time for the long road*. A participant's looking time for a road was the sum of all the times that participant looked at that road while the car was under the occluder in milliseconds.

The variable looking time was computed by subtracting the looking time for the short road from the looking time for the long road and dividing that number by the total looking time.

$$\textit{Looking time} = \frac{\textit{(looking time long road – looking time short road)}}{\textit{(looking time long road + looking time short road)}}$$

Results

Anticipation

We expected that participants would anticipate first to the short road. However, looking at Table 1 and Figure 5, the opposite seems to be true. Also, there were 6 participants that did not anticipate to any of the roads while the car was under the occluder (see Table 1). To see if the difference in anticipation between the short road and the long road was significant, a *t*-test for a single sample was done, comparing anticipation to the value 0. For this and every other analysis in this paper, we adopted the standard alpha level of .05. There was no significant difference in anticipation between the short (N = 7) and the long (N = 14) road ($t(26) = .1.568$, $p = .129$), indicating that participants did not anticipate more the short road or to the long road. However, this difference might be significant with more participants.

		Anticipation			
		Short	None	Long	Total
Condition	Race	5	2	6	13
	Control	2	4	8	14
	Total	7	6	14	27

Table 1: Participant count for anticipation and condition

We expected that participants would anticipate more to the short road in the race condition than in the control condition. However, looking at Table 1 and Figure 5, it seems that participants anticipated equally as much to short as to the long road in the race condition, and more to the long road in the control condition. To see if this difference in conditions was significant, a chi-square test was done for anticipation and condition. The percentage of participants that looked first at a certain road did not differ by condition, $\chi^2(2, N = 27) =$

2.204, p = .332, indicating that participants did not anticipate differently in the race condition than in the control condition.

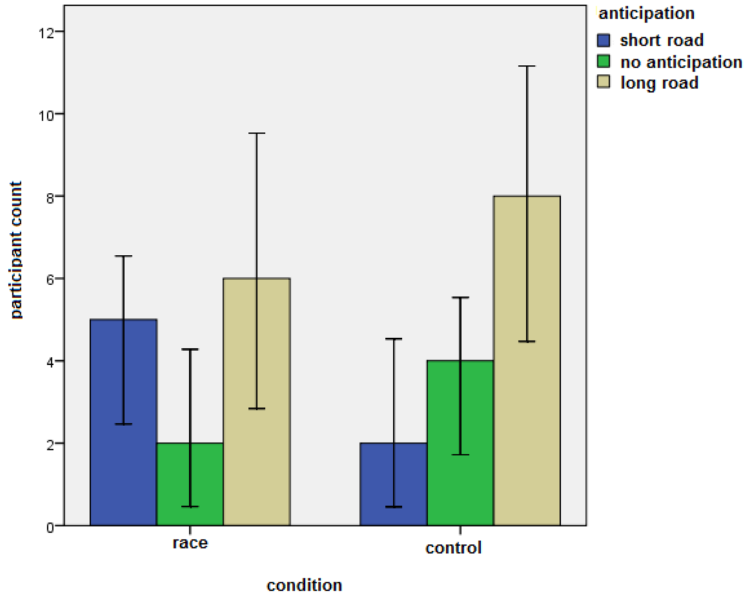

Figure 5: Anticipation for the short road and for the long road divided over the race condition and the control condition

Looking time

We expected that participants would look longer at the short road than at the long road. However, just like for anticipation, Table 2 and Figure 6 suggest that the opposite is true. A paired-samples *t*-test was conducted to compare the looking time short road and for the long road. There was a significant difference between the looking times for the short (M = 300.4, SD = 488.2) and the long (M = 646.4, SD = 582.2) road (t = -2.726, p = .011), indicating that participants looked longer at the long road than at the short road.

		Mean looking time (ms)		
		Short	Long	Total
Condition	Race	361.5	754.9	1116.5
	Control	243.6	545.6	789.1
	Total	300.4	646.4	946.7

Table 2: Mean looking time for roads over conditions

We expected that participants would look longer at the long road in the race condition than in the control condition. However, looking at Table 2 and Figure 6, it seems that the difference in looking time between the roads is equal for both conditions. An independent-samples *t*-test was conducted to compare looking time in the race and the control condition. There was no significant difference in the scores for the race (M=.4, SD=.5) and the control (M=.4, SD=.6) conditions ($t(25) = .260$, $p = .797$). These results indicate that anticipation is not dependent on the race context.

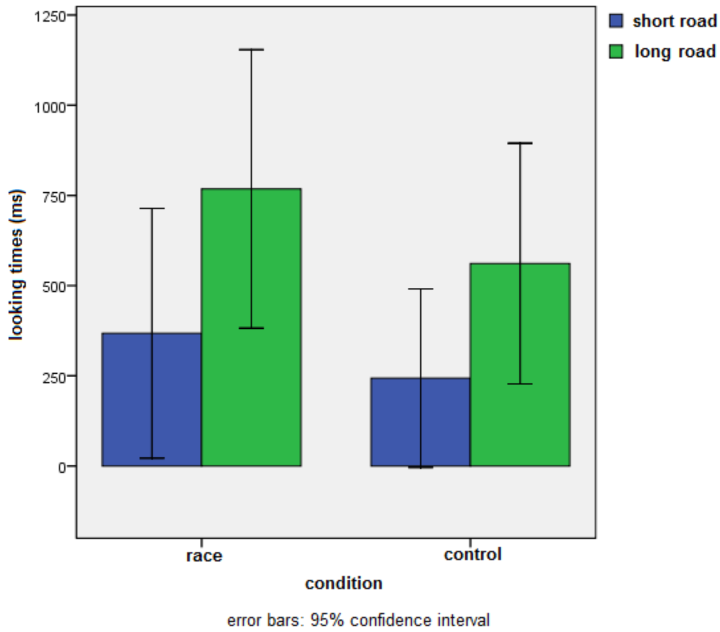

Figure 6: Looking time for the short road and for the long road divided over the race condition and the control condition

Figure 7 suggests a difference in looking time between the long right and the long left stimuli. To see if this difference was significant, a *t*-test for independent-samples was conducted to compare looking time for the long right and the long left stimuli. There was no significant difference in looking time between the long right (M = .6, SD = .4) and the long left (M = .2, SD = .6) stimuli ($t\ (25) = 1.786$, $p = .086$), but because the effect was marginally significant, the data for the long right (see Figure 8) and long left stimuli (see Figure 9) were analyzed separately.

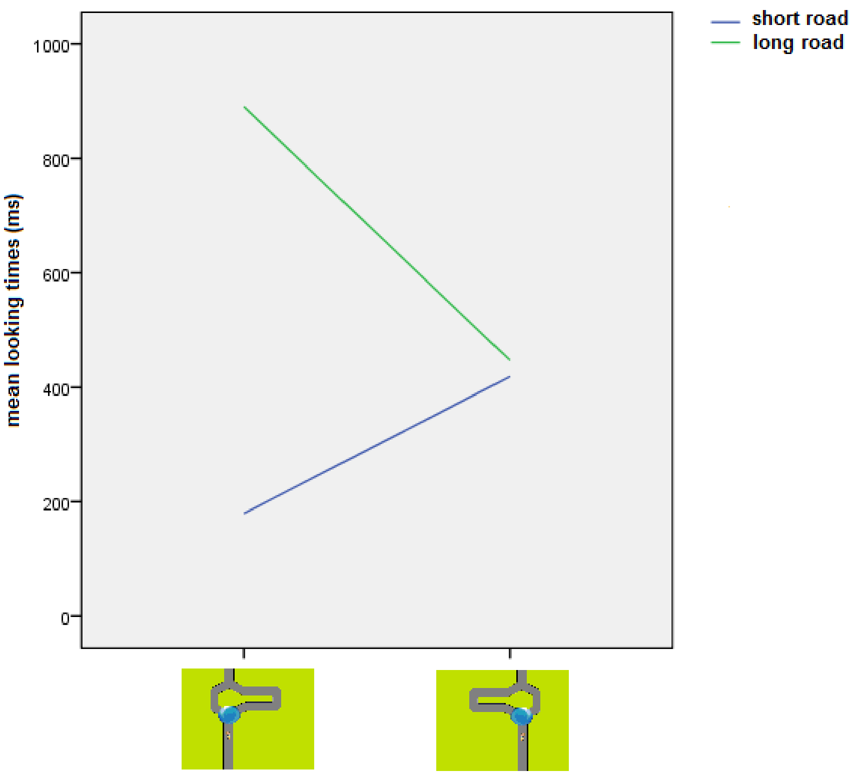

Figure 7: The effect of mirroring on looking times

Long right stimuli. Figure 8 suggests that participants who received the long right stimuli looked longer at the long road than at the short road. To see if this difference was significant, a *t*-test for independent-samples was done. For the long right stimuli, there was a significant difference between the looking time for the short (M = 179.4, SD = 247.3) and the long (M = 860.8, SD = 617.0) road ($t(12) = -4.680$, $p < .001$), indicating that participants who received the long right stimuli looked longer at the long road than at the short road.

Looking at figure 8, it seems that the difference in looking time between the roads is equal for both conditions. A *t*-test for independent-samples was conducted to compare looking time in the race and the control condition for the long right stimuli. There was no significant difference in the scores for the race (M=.6, SD=.4) and the control (M=.6, SD=.5) conditions ($t(11) = .123$, $p = .904$). These results indicate that anticipation of participants who received the long right stimuli is not dependent on the race context.

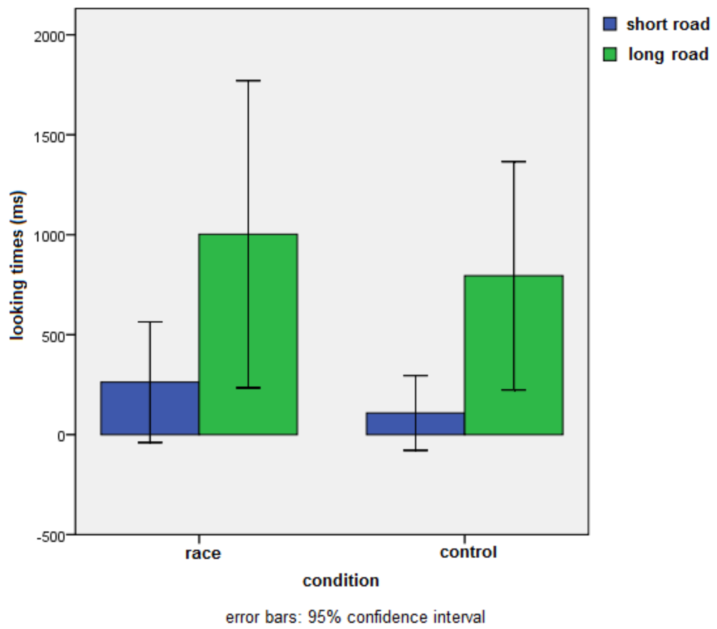

Figure 8: Looking times for the long right stimuli

Long left stimuli. If one looks at figure 9, there seems to be no difference between the looking time for the short road and the long road for participants who received the long left stimuli. A *t*-test for paired-samples was conducted to compare the looking time for the short and the long road. There was no significant difference in looking time between the short (M = 412.7, SD = 626.2) and the long (M = 447.3, SD = 487.7) road ($t(13) = -.205$, $p = .841$). These results indicate that participants who received the long left stimuli did not anticipate more to either of the roads.

Figure 9 suggests that in the race condition the looking time for the long road is longer and in the control condition the looking time for the short road is longer. A *t*-test for independent samples was conducted to compare looking time in the race and the control condition for the long left stimuli. There was no significant difference in the scores for the race (M=.3, SD=.5) and the control (M=.1, SD=.7) conditions ($t(12) = .317$, $p = .757$). These results indicated that anticipation of participants who received the long left stimuli is not dependent on the race context.

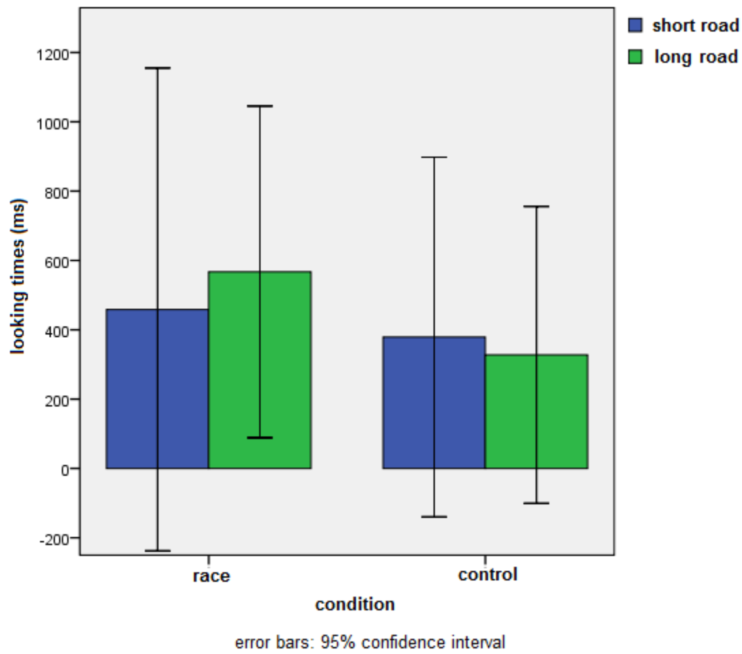

Figure 9: Looking times for the long left stimuli

When looking at the gaze data, it became clear that the period that the car was under the critical occluder was really long (3000 ms). Participants shifted their gaze from the short to the long path a lot. For this reason, we separated the time when the car was under the occluder into periods of 200 ms (see Figure 10). We looked at the 1000 ms period around the time when the car should have reappeared from under the occluder (400 ms to 1400 ms from the time when the car disappeared behind the occluder.

Table 3 and Figure 10 suggest that participants looked longer at the long road than at the short road. A paired-samples *t*-test was conducted to compare the looking time short road and for the long road. There was a significant difference between the looking times for the short (M = 125.6, SD = 196.5) and the long (M = 292.6, SD = 323.1) road (t = -2.285, p = .031), indicating that participants looked longer at the long road than at the short road in the period when the car should have reappeared from under the occluder.

Looking at Table 3, it seems that the difference in looking time between the roads is equal for both conditions. An independent-samples *t*-test was conducted to compare looking time in the race and the control condition for the period when the car should have reappeared from under the occluder. There was no significant difference in the scores for the race (M=.2, SD=.8) and the control (M=.4, SD=.7) conditions ($t(25) = -.511$, $p = .614$). These results indicate that anticipation is not dependent on the race context in the period when the car should have reappeared from the occluder.

		Mean looking time (ms)		
		Short	Long	Total
Condition	Race	150.4	316.2	466.6
	Control	102.6	270.6	373.1
	Total	125.6	292.6	418.1

Table 3: Mean looking time for roads over conditions for the period when the car should have reappeared from under the occluder

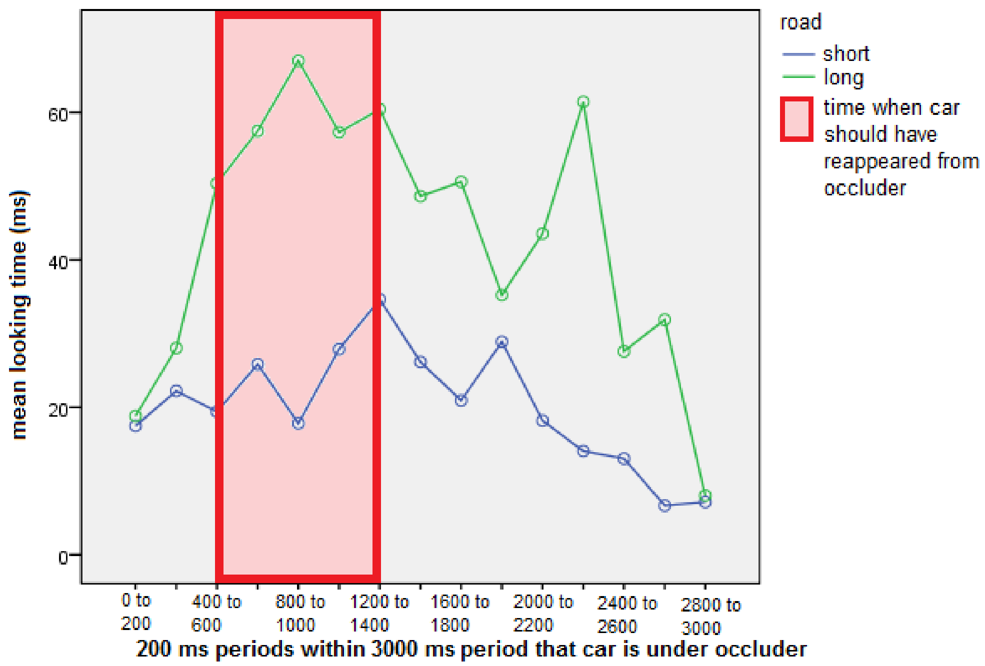

Figure 10: Mean looking times for 200 ms periods within 3000 ms period that car is under occluder

Discussion

In this study, we tried to answer two questions: (1) Do people predict other agents' actions assuming that they behave efficiently? (2) Is this dependent on the context?

The results of Gergely et al. (1995) and Csibra et al. (1999) were interpreted as a principle of rationality by Gergely and Csibra (2003) and as a frequency effect by Paulus et al. (in press). When Paulus et al. (in press) tried to disentangle rationality and frequency, their results suggested that adults relied on rationality information after their predictions based on frequency information were not met (Paulus et al., in press). This experiment was run to see if people expect other agents to behave rationally when frequency information was unavailable and to see if this was dependent on the context.

We expected that participants would anticipate more to the short road than to the long road, because other studies suggested rationality (Gergely et al., 1995; Csibra et al., 1999; Paulus et al., in press). We found no evidence for this prediction. The opposite was found. The looking time was significantly higher for the long road than for the short road. Further, we expected that participants would anticipate more to the short road in the race condition than in the control condition, because in a race context, a car would have to hurry more than in an everyday context. We found no difference between the two conditions.

One might argue that there is no effect of condition on anticipation or looking time, because our attempt at creating a sense of hurry failed. In hindsight, we could have asked the participants whether they felt that the car had to hurry. In this way, we would have known if participants in the race condition felt differently about this than participants in the control condition.

The fact that overall people looked longer at the longer road suggests that, contrary to our expectation and findings reported by others (Csibra et al. (1995), Gergely et al. (1999)), people do not expect other agents to behave rationally. Alternatively, our findings may have been artifacts or confounds in the design. For example, participants might still be exploring the stimuli with their gaze while the car goes under the occluder. Exploring the long road would take more time than exploring the short road. Participants might look longer at the long road, simply because there is more to look at.

However, our AOIs have the same size for both roads and are positioned in the occluder area. Additionally, when looking at the gaze replays of the participants, we found that most participants’ gazes stayed in the vicinity of the occluder. Their eyes did not wander to the rest of the long road. This suggests that there is anticipation at play, only not to the short road, as we expected, but to the long road. Possible explanations are that participants do not see efficiency as rationality or that they do not predict other agents’ to behave rationally.

However, we found that people only look longer at the long road when it is on the right side of the screen. When the long road is on the left side of the screen, there is no significant difference in looking times between the short road and the long road. Only anticipation to the long road cannot explain the difference between the long right and the long left stimuli.

However, this might be explained by a known effect that people have a preference for the right side of stimuli (Cohen, 1972; Chokron & De Agostini, 2000). This may explain the difference between the long right and the long left stimuli. Also, the fact that Dutch people drive on the right side of the road might explain why people expected the car to drive on the right road more. It seems there are two effects in play: a preference for the right side and anticipation to the longer road. This explanation might be tested in future studies. For example, the same study might be repeated with one road above and another road below. However, there might also be an above or below preference.

The results of the present study shed a different light on the results of the experiment of Paulus et al. (in press). As you may recall, they found that infants and adults looked at the long road in the first trial. They explained this as a frequency effect, because they had shown the cow taking the long road several times before in the habituation phase. However, if, in our single trial experiment, people expect other agents to take the less efficient road, then it might be that their results are not explained by a frequency effect, but by a general first trial preference for the longer road. There might be an anti-rational or an anti-efficient bias.

If people do not base their expectations about other agents' actions on an assumption of rationality, maybe robots should not do so either. One might argue that it does not matter whether robots do things in the same way as humans. However, a robot that is wired like a human might communicate in a more natural way with a human. Since, for a social robot, communication is an essential feature, it makes sense to build the robot in the same way as a human works.

Now, looking back at our restaurant example, imagine maybe one time it will be possible to go to a robot restaurant and feel like these robots are attuned to your eating behavior in much the same way as a good human staff can be.

References

Baker, C. L., Saxe, R., and Tenenbaum, J. B. (2009). Action understanding as inverse planning. *Cognition, 113,* 329-349.

Baker, C. L., Tenenbaum, J. B., and Saxe, R. R. (2006). Bayesian models of human action understanding. *Advances in neural information processing systems, 18,* 99-106.

Baker, C. L., Tenenbaum, J. B., and Saxe, R. R. (2007). Goal inference as inverse planning. In *Proceedings of the twentyninth annual conference of the Cognitive Science Society.*

Blokpoel, M., Kwisthout, J., van der Weide, T. P., and van Rooij, I. (2010). How action understanding can be rational, Bayesian and tractable. *Manuscript under review for Proceedings of CogSci2010.*

Breazeal, C. (2003). Emotion and sociable humanoid robots. *International Journal of Human-Computer Studies, 59,* 119-155.

Chokron, S., and De Agostini, M. (2000). Reading habits influence aesthetic preference. *Cognitive Brain Research, 10,* 45-49.

Cohen, L. B. (1972). Attention-getting and attention-holding processes of infant visual preferences. *Child Development, 43,* 869-879.

Csibra, G., Gergely, G., Bíró, S., Koós, O., and Brockbank, M. (1999). Goal attribution without agency cues: The perceptions of 'pure reason' in infancy. *Cognition, 72,* 237-267.

Dautenhahn, K., and Billard, A. (1999). Bringing up robots or-the psychology of socially intelligent robots: From theory to implementation. *Proceedings of the Third International Conference on Autonomous Agents.*

Forlizzi, J., and DiSalvo, C. (2006). Service robots in the domestic environment: a study of the Roomba vacuum in the home. *Proceedings of HRI06,* 258-265.

Fujita, M. (2001). AIBO: Toward the Era of Digital Creatures. *The International Journal of Robotics Research, 20(10),* 781.

Gergely, G., and Csibra, G. (2003). Teleological reasoning in infancy: the naïve theory of rational action. *TRENDS in Cognitive Sciences, 7*(7), 287-292.

Gergely, G., Nádasdy, Z., Csibra, G., and Bíró, S. (1995). Taking the intentional stance at 12 months of age. *Cognition, 56,* 165-193.

Heerink, M., Kröse, B., Evers, V., and Wielinga, B. (2006). Studying the acceptance of a robotic agent by elderly users. *International Journal of Assistive Robotics and Mechatronics, 7(3),* 33-43.

Paulus, M., Hunnius, S., van Wijngaarden, C., Vrins, S., van Rooij, I., and Bekkering, H. (in press). The role of frequency information and teleological reasoning in infants' and adults' action prediction. *Developmental Psychology*.

Prassler, E. Ritter, A. Schaeffer, C., and Fiorini, P. (2000). A short history of cleaning robots. *Autonomous Robots, V9,* 211-226.

Reeves, B., and Nass, C. (1996). The Media Equation. CSLI Publication, Stanford, California.

Van Breemen, A., Yan, X., and Meerbeek, B. (2005). iCat: an animated user-interface robot with personality. *Proceedings of the Fourth International Joint Conference on Autonomous Agents and Multiagent Systems,* 143-144.

Wada, K., Shibata, T., Saito, T., and Tanie, K. (2003) Effects of robot assisted activity to elderly people who stay at a health service facility for the aged. *Proceedings of the 2003 IEEE/RSJ Intl Conference of Intelligent Robots and Systems.*