

How a cow catches a sheep

*Is action anticipation in infants based on
rationality or frequency?*

Written by Carolien van Wijngaarden

0622729

03-07-09

Bachelor internship and thesis

Artificial Intelligence and Cognitive Psychology

Supervisors dr. Iris van Rooij, drs. Markus Paulus and dr. Sabine Hunnius

Abstract

The aim of social robotics is to make robots that are integrated in our daily life and cooperate with humans. For efficient cooperation, the understanding of actions of other agents is important. In Artificial Intelligence the general assumption is that humans will perform the actions that are the most rational way to achieve a goal. Infant studies about action anticipation performed by Gergely, Nádasdy, Csibra and Bíró, (1995) and Csibra, Gergely, Bíró, Koós and Brockbank (1999) are often cited as support for the viability of that assumption. The present study, however, investigates the possibility that action anticipation is based on the frequency of an action instead of rationality. To test this hypothesis we performed an experiment using a habituation paradigm in which we measured the looking time as well as the anticipation of 9-month-old infants when they observed an agent performing one out of two possible actions. We manipulated the actions insofar as one of the actions was the more frequent but also more inefficient one, whereas the other was the more efficient but also more infrequent one. The anticipation measurements showed evidence for the frequency hypothesis, whereas the looking times provided no evidence for either the frequency hypothesis or the rationality theory. Therefore, it could be interesting to see how action models in Artificial Intelligence based on frequency will perform in comparison with or in cooperation with existing models based on rationality.

INTRODUCTION

A few days ago I walked into a store at the station to get some food. My favourite sandwich was placed on the top shelf and there was exactly one left. I reached for the sandwich but it was too high for me to reach it. I looked around but there was nobody nearby whom I could ask for help. Frustrated, I took a different sandwich and hurried to the cash register. Having lost too much time buying my food, I ran to the train. I reached the door, but I could not press the button to open the doors fast enough. Before I knew it the train was gone and I was eating a bad sandwich. Pfff... I hate Monday mornings.

Could my Monday mornings be improved? Looking closely at the events in my Monday morning, I realized something: none of my problems would have arisen if there had been some help around understanding my actions. For example, if there was somebody in the store observing me that Monday, he could see me reaching for the sandwich and help me take it, before I even realized I could not get it myself. Of course, there cannot always be persons around me observing my actions the whole time. Such job is too boring and too tiring for a human, but would it not be a perfect job for a robot? For a robot with an understanding of actions, integrated in our everyday life?

A research field in Artificial Intelligence dealing with robots meant to function in the everyday life of a general public is the field of social robotics. Social robots are autonomous robots that display social behaviours and comply with social rules in interacting and communicating with humans or other autonomous agents (Sabanovic, Michalowski & Simmons, 2006). For efficient interaction between social robots and humans, it could be useful for robots to have the ability to predict and understand the actions humans are likely to make in a given setting.

The “Monday morning” problem seems insignificant, but if you imagine the person in the story to be in a wheelchair, it becomes clearer that for him or her, the action problems in the Monday morning story are more important and occur frequently. Also in other situations social robots with an understanding of actions could be very useful, for example a social robot could assist a medical doctor during surgery. If the social robot could predict the actions of a doctor, it could hand him the instruments in the right order and change this order if the actions of the surgeon do not meet his expectations.

How can the ability to understand actions be implemented in a robot? One approach to this question is to take the only example of a system we know that has this ability as a source of inspiration: humans. Over the years, robotics researchers have often tried to reach human-level intelligence in robots by trying to immediately program “grown-up” human intelligence. Although this approach has had its notable successes for relatively abstract or rule-based tasks (e.g., chess), it has proven difficult to achieve everyday common sense behaviour in artificial systems (Haselager, 1997; Pylyshyn, 1987). One approach that may overcome this limitation is based on the idea that biological systems, like us, do not come into the world equipped with all their basic cognitive abilities; instead these abilities develop over time. Humans start as infants with some initial conditions and by responding to environmental changes they continuously improve their capabilities and even develop abilities which were non-existent before. Analogously, a capability may be developed in robots by first programming infant-level intelligence and then let this robot “grow up” either through experience and imitation or explicit programming of different developmental stages (Turing, 1950; Bakker & Kuniyoshi, 1996; Demiris & Meltzoff, 2008). In either case, a natural starting point is to study the relevant ability in infants.

There is a considerable amount of evidence that infants can attribute goals to actions in a way that allows them to predict the future course of the actions (Csibra, 2003; Hauf & Prinz, 2005; Woodward, Sommerville & Guajardo, 2001). Two ways in which goal attribution allows predicting future events are *goal prediction* and *action anticipation* (Csibra & Gergely, 2007). Goal prediction is predicting the function (or goal state) of an action when the action is not yet finished (an ‘action-to-goal’ inference). For example, someone watching me reach for the sandwich could predict that the goal object of my action is the sandwich. Action anticipation is making anticipations about what actions would achieve a goal (a ‘goal-to-action’ inference). For example, if I see a man running to a train (the goal object), I could predict his next actions will be running further to the train, opening the door of the train and getting in. In this thesis I will concentrate on action anticipation, because the big advantage of action anticipation is that if the predictions about future actions are not confirmed we can revise our whole action interpretation. For example if I see the running man passing the train and run towards a building, my predicted actions of opening the door of train and getting in are not confirmed. Therefore I would have to change the predicted goal of the running man and my anticipations of the actions he will perform

to reach this new goal. I could change my prediction of the goal from ‘going to the train’ in ‘going to the toilet’ and my anticipation of future actions would become going to the toilet and paying the toilet cleaner.

How is action anticipation achieved in infants? One prominent hypothesis about infants’ abilities to predict actions of others is based on the ‘principle of rational action’ (Gergely & Csibra, 2003). This principle states that we understand others actions by assuming that actions are ‘rational’ (e.g., as effective and efficient as possible). Two forms of understanding can be distinguished: (1) when observing actions, infants predict those goals that are most efficiently achieved by the observed actions, and (2) given a particular goal, infants predict the most efficient action to reach the goal.

An often-cited study by Gergely, Nádasdy, Csibra and Bíró, (1995) purports to provide evidence that infants evaluate the actions of others according to the ‘principle of rational action’. In their experiment, 12-month-old infants watch a movie in which a ball (suggested to be an intentional agent) goes to another ball and makes a jumping movement to get there. In the experimental condition there is an obstacle on the pathway between the two balls. In the control condition there is no obstacle between the two balls (see Figure 1). After a habituation period infants in both groups see the ‘old’ and the ‘new’ action. In the old action event the small ball makes exactly the same jumping movement as in the habituation period even though there is no obstacle between the two balls. In the new action event, the ball makes a straight-line movement along the floor to the other ball. Gergely et al. found that infants in the experimental condition looked longer to the old than to the new action. Their interpretation of this finding is that infants are more surprised by the old action than the new action, because the old action is less rational than the new action. According to Gergely et al. the result of this experiment is evidence for the hypothesis that infants reason rationally about actions of other agents. Precisely the same experiment was repeated with 9-months-olds by Csibra, Gergely, Bíró, Koós and Brockbank (1999) and the same results were found.

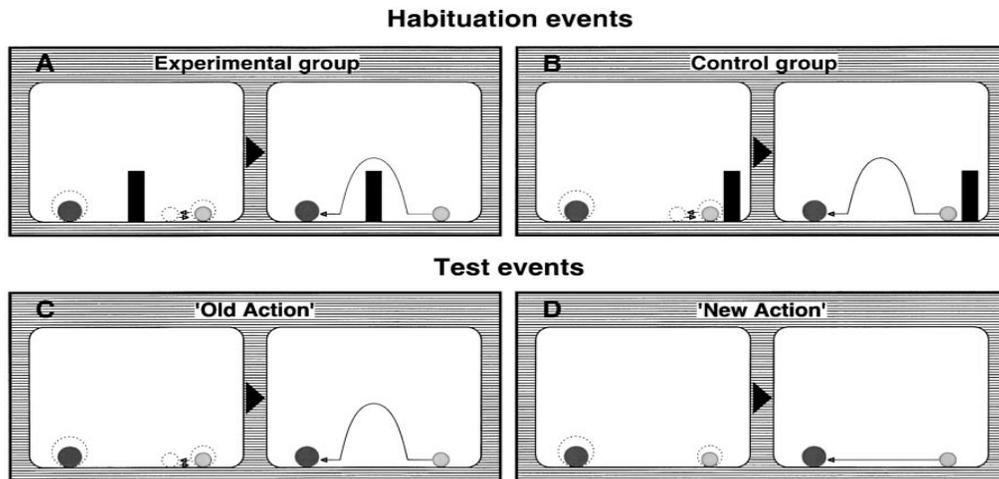


Figure 1: Illustrations of the experimental stimuli from the paper of Csibra et al. (1999).

Many computational models of action understanding in cognitive neuroscience and AI are similarly based on such a ‘principle of rational action’ in humans (Baker, Tenenbaum, & Saxe, 2007; Oztop, Wolpert & Kawato, 2004; Mao & Gratch, 2004; Youn & Oh, 2008; Verma & Rao, 2006). These studies try to simulate goal attribution and action anticipation based on the principle of rationality by also making representations for beliefs, intentions and desires (Davidson, 2001; Youn & Oh, 2008). The infant studies of Gergely et al. (1995) and Csibra et al. (1999) are often cited as support for the viability of the assumption that humans base their prediction of action of others on rational reasoning. Others have also claimed to have found evidence for the principle of rationality in infants using very different experimental set-ups (Schwier, van Maanen, Carpenter, & Tomasello, 2006; Csibra, 2003; Gergely, Bekkering, & Király, 2002). Such studies, however, have been shown to allow for more parsimonious explanations (Sirois & Jackson 2007; Paulus, Hunnius, Vissers, & Bekkering, 2009). Paulus et al., for example, showed that infants’ imitation behaviour is caused by motor resonance (i.e. the mapping of the others’ actions onto one’s own motor repertoire) instead of a rational evaluation of others’ actions. One may wonder, thus, if similarly parsimonious explanations are causing the Gergely et al. (1995) and Csibra et al. (1999) finding.

Looking at the two test events of the experiment of Csibra et al. (1999) the new action is the most rational action of the two, but it is also the most frequently seen action in daily life. Arguably infants see movements that follow a surface more often, than ‘jumping’ movements, so that the start frequency of the new action is

higher than that of the old action. Therefore, the infants' expectations as demonstrated in the experiment of Csibra et al. could possibly also be based on statistical learning (henceforth: frequency hypothesis) instead of rationality. Given the dominance of rational models of action understanding in AI, and cognitive science in general, addressing this issue is important, because if the existing models do not simulate the way humans deal with actions, the models may never reach human level intelligence, the central objective of Artificial Intelligence.

To test the frequency hypothesis we designed an experiment that pits the frequency hypothesis against the rational theory. In our study an agent could walk on two pathways to reach a goal. The rationality of the action 'taking a pathway' was manipulated by making one pathway short and the other one long. The frequency of the action was manipulated by habituating the infants with only the presentation of the agent taking the long pathway to reach the goal. Following either the long or the short pathway led to a movement following a surface, thus removing the possible difference in the start frequency of both actions noticed in the experiment of Csibra et al. (1999). In the habituation phase, taking the long pathway was the only possibility to reach the goal, because there was a gap in the short pathway. In the test phase both pathways were intact and infants were shown two actions: the agent taking the short pathway (the most rational but least frequently shown pathway) and the agent taking the long pathway (the least rational but most frequently shown action).

In Csibra et al. (1999) the visual habituation paradigm (cf. Spelke, 1985) was used to test the hypothesis that infants judge actions based on the principle of rationality. Measuring infants' looking time is an indirect way to measure expectations about actions. However, with current technology there is a more direct technique available for measuring action anticipation namely eye-tracking (Falck-Ytter, Gredebäck & von Hofsten, 2006). With the eye-tracking technology, we can measure where an infant is looking during the presentation of the stimuli. Action anticipation could be measured with an eye tracker by letting an agent walk under an occluder at the moment the agent has to choose between two actions and measure where the infants look during the period the agent is under the occluder. Therefore, in this study we measure both looking time and anticipation to test our hypothesis that action anticipation in infants could possibly be based on frequency instead of rationality.

If infants look longer to movies in which the agent takes the long pathway (irrational action) and at the same time show anticipatory looks to the short pathway this is evidence in favour of the rationality theory. However, if our frequency hypothesis is true we expect the infants to look longer to the movies in which the agent takes the short pathway, and additionally we expect the infants to anticipate to the long pathway.

METHOD

Participants

Thirty-eight 9-month-old infants (14 girls, 24 boys) participated in this experiment. The data of nineteen infants (5 girls, 14 boys) were used for analysis; nineteen infants were excluded from the analysis. Infants were excluded from analysis for the following reasons: if the infant became fussy or started crying, if the parents interfered with the infant or if the infant was not habituated and thus had not seen the habituation movie enough times to recognize the movie as something they had seen before. The infants who did not look at all to the important parts of the test trials were also excluded from the analysis. This last point is explained in more detail in the procedure section.

The remaining infants were on average 9 months and 21 days old (range 9 months, 16 days to 9 months, 27 days). The parents of the infants were contacted through information from public birth records. The infants were healthy and had no pre- or perinatal complications. All the parents agreed to sign a consent form which stated that we could tape the infant's face and use the recording to code the infant's looking times. As compensation for participating in the experiment, the parents could either choose a baby book or a monetary compensation.

Apparatus

The experiment was programmed in Presentation 11.07 (Neurobehavioral systems USA). Infants' eye gaze was recorded using a corneal reflection eye-tracker (Tobii 1750, Tobii Technology, Stockholm, Sweden), integrated in a 17" TFT flat-screen monitor. For the calibration of the Tobii the Clear View software (Tobii Technology AB) was used.

If an infant was moving his or her head too much to one the side, there was a chance the eyes of the infant were out of sight for the Tobii, leading to errors in

looking time. Two digital cameras (Sony Handycam DCR-SR190E) were used to check if the child stayed in reach of the eye-tracker; one recorded the face of the infant, the other one the screen.

Stimuli

The stimuli consisted of three different premovies, a habituation movie and two types of test movies. The stimuli (short movies of 1200 by 1024 pixels (px)) were inspired by the stimuli of Csibra et al. (1999) and were made in Adobe Image Ready 7.0. In the movies an agent (a cow) could choose between two pathways (short and long) to reach another agent (a sheep). Like in the experiment of Csibra et al. the movies started with an (non verbal) action sequence, which can be seen as a communication between the cow and the sheep. After the communication period, the cow started to walk to follow the sheep, but before the cow chose a pathway he disappeared under an occluder overlaying the crossway between the short and the long pathway. The time the cow was under the occluder was long enough to give the infants the opportunity to anticipate at which pathway the cow would appear, but it was short enough give the idea that the cow was still walking while it was under the occluder. After this period the cow chose a pathway and walked off screen towards the sheep. Below I will explain all stimuli in more detail.

The movies in the habituation period started with a 0.2 seconds (s) resting period in which the sheep was standing at the right and the cow was standing at the left part of the screen facing the sheep (see Figure 2A). The cow and the sheep were both standing on a pathway, but in between them the pathway splits in two. A long pathway was leading to the sheep and a short pathway would lead to sheep if the gap in the middle of the pathway would not be there. A blue transparent oval occluder overlaid the left crossway between the short and the long pathway.

After the resting period the sheep wiggled (0.7s), waited for 0.5s and wiggled away to the right until it was off screen (0.8s). When the sheep was gone, the cow waited (0.5s) and than also wiggled for 0.7s. The wiggling sequence of the sheep and cow could be seen as a form of communication.

When the communication period was ended the transparent blue oval occluder which was covering the crossway between the long and the short pathway was gradually turning non-transparent in 0.8s. The cow walked (0.8s) under the non-transparent oval and stayed there for 1.5s. The cow appeared at the long pathway and

walked (4.5s) on this pathway, to the place where the sheep stood in the beginning of the movie, and walked off screen. The movie ended with a black screen for 1s. The whole movie took 12s.

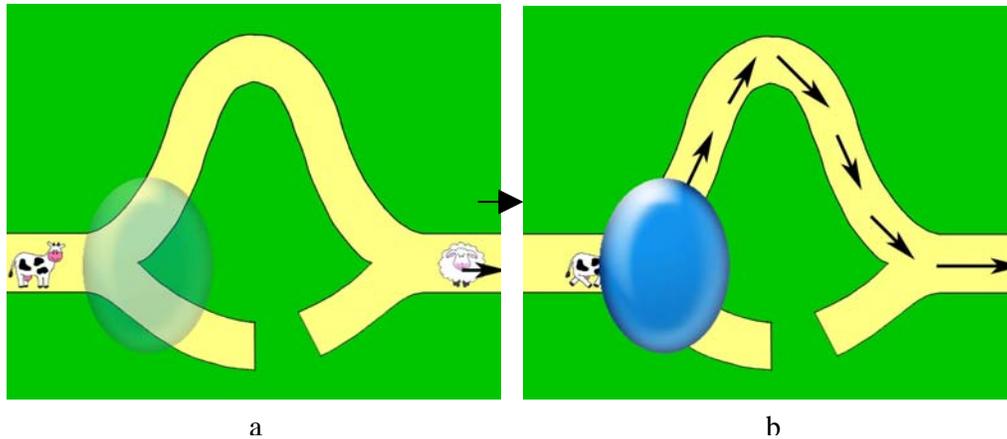


Figure 2: Images of the habituation movie showing the long pathway at the top and the short pathway at the bottom. (a) The beginning of the movie where the sheep communicated with the cow and then walked off screen. (b) After the communication period the cow walked under the (now non-transparent) occluder, took the long pathway and walked off screen.

Just like in the experiment of Csibra et al. (1999) the test phase consisted of two test trials. In the movies of both test trials the gap in the middle of the short pathway is gone. One test trial (old action) showed the cow taking the long pathway and the other (new action) one showed the cow taking the short pathway. Both test trials started with a 1s resting period.

The old action movie (see Figure 3) was exactly the same as the movie of the habituation period except for the difference that the gap in the short pathway was gone. This means that the short pathway also led to the sheep.

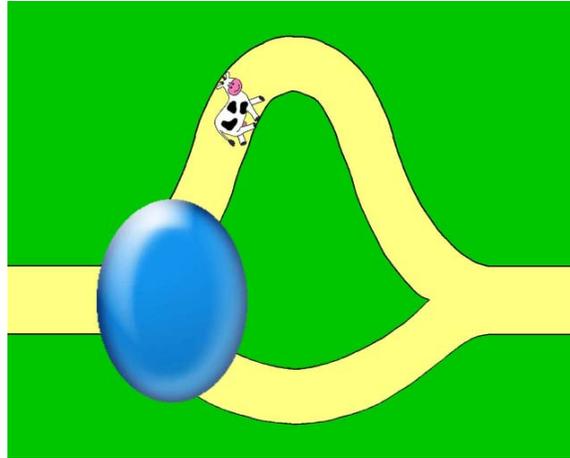


Figure 3: Image of the old action movie. The gap in the short pathway is gone, but the cow still took the long pathway and walked off screen.

The new action movie (see Figure 4) was exactly the same as the old action movie for the first 6.5s (until the cow was coming from under the oval). Instead of appearing at the long pathway the cow appeared at the short pathway, walked for 3.4s to the place where the sheep stood in the beginning of the movie, and walked off screen. This movie too ended with a black screen for 1s. This movie is 1.1s shorter than the old action movie and takes 10.9s to finish (cf. Csibra et al., 1999).

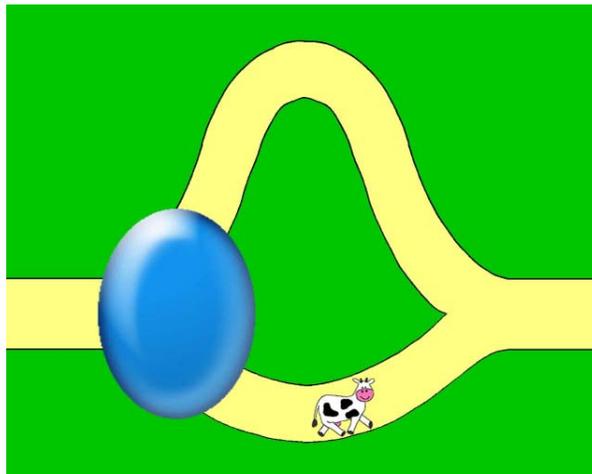


Figure 4: Image of the new action movie. The cow took the short pathway and walked off screen.

To ensure that infants were familiar with the setting of the movies, a number of premovies were used which demonstrated what paths there were, which animal was involved and what the occluder was. The premovies (Figure 5) consisted of a straight pathway starting at the right side and ending just over the middle of the screen at the

left side. After waiting for 0.2s the cow walked from the right side to the end of the pathway, turned around and walked back. The total duration of the premovies was 5.2s. The difference between the three premovies was the state of the occluder. The first premovie had no blue oval in the middle, the second one had a 59% transparent blue oval in the middle and the last premovie had a non-transparent blue oval in the middle. In the first two premovies the cow could be seen during the whole movie. In the last premovie the cow disappeared and appeared twice behind the blue oval during the movie.

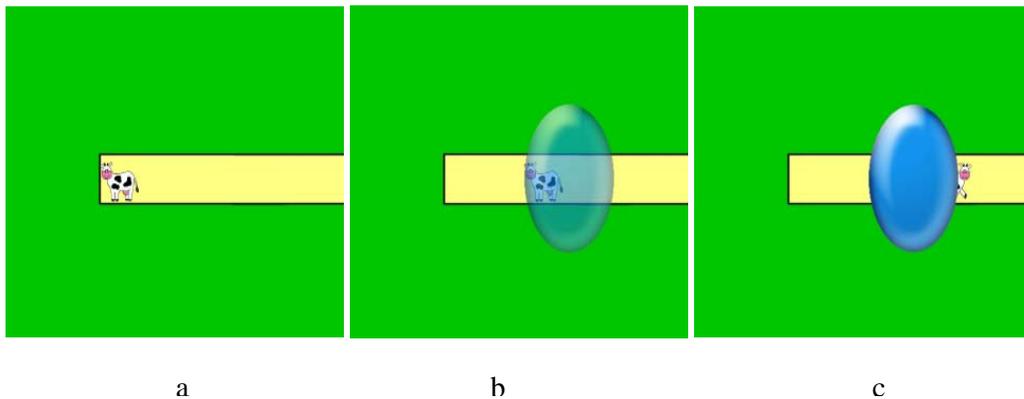


Figure 5: Images of the three different premovies. The actions of the cow were equal in all movies, but the movies differed in the transparency of the occluder. (a) An image from the first premovie. This premovie showed no occluder. (b) An image from the second premovie, in which the occluder was transparent. (c) An image from the last premovie, in which the occluder was non-transparent.

Design

The experiment consisted of a within-subjects design with Action Type (new action, old action) as independent variable and Looking Time and Anticipation as dependent variables. We counterbalanced the order of the independent variable Action Type and the order of the long and short pathway (at top, at bottom). This can be seen in table 1.

		Order of Action Type		Total
		Old action first	Short action first	
Placement of the pathways	Long pathway top, short pathway bottom.	4 infants	5 infants	9 infants
	Long pathway bottom, short pathway top.	5 infants	5 infants	10 infants
	Total	9 infants	10 infants	19 infants

Table 1: This table shows how Action Type and the order of both pathways were counterbalanced. For each counterbalancing group the number of infants in that group is shown.

Procedure

The experiment was carried out in a room with lightproof curtains and two white curtains separating the test materials from the other materials in the room. Before the experiment started, the parents were instructed not to interact with their infant during the experiment in order to minimize the distraction of the infant. The infant was placed in front of the Tobii eye-tracking system in an infant seat, on the parent's lap, at a viewing distance of approximately 60 cm. The monitor was connected to a movable and adjustable arm, so that the screen could always be put in the correct position. The screen was adjusted so that it was parallel to the face of the infant. After adjusting the screen to the correct position the experimenter turned on both cameras making sure the face of the infant and the screen were in focus. The experimenter turned off all the lights, so the infant's attention could not easily be drawn to other things in the room.

Prior to testing, the gaze of each infant was calibrated by using a 9 point calibration procedure. For every calibration point in a screen-wide 3x3 grid, a white expanding-contracting circle on a black square background appeared. When the infant looked at the figure, the experimenter pressed a button on a keyboard resulting in the appearance of the figure at another point. This process was repeated until all nine points were calibrated correctly.

During the experiment the Tobii eye-tracker recorded the gaze of the infant. The experiment started with an attention-getter of 10s. In these 10 seconds the experimenter hid herself behind a curtain where she, during the experiment, followed the course of the experiment at a TV screen. Then the premovies were started. The premovie without blue oval was shown once, the premovie with a transparent blue oval in the middle was shown twice and the last premovie with the non-transparent blue oval was also shown twice.

After the premovies were shown the habituation period started. At the beginning of each trial the attention of the infant was directed to the screen by showing an attention-getter with an accompanying sound. When the infant looked at the screen the trial was started and the habituation movie was shown repeatedly. A trial ended when the infant looked away for longer than 2s or when the trial reached its maximum duration of 1 minute. The computer averaged the looking time of the first three trials and compared this value on-line with the last three trials. The habituation criterion was: the average looking time of the last three trials had to be less than 50% of the average looking time of the first three trials and this criterion had to be met twice in a row. Thus the minimal number of habituation trials was seven. The maximal number of habituation trials was set to 15.

In the original Csibra et al. (1999) experiment a 30s break was introduced after the habituation criterion was reached in which the parent and the infant turned away from the screen. This was not possible in the current experimental setup, because by moving the infant we could lose the calibration data. Therefore, in this experiment a 30s break was introduced by showing the infants an unrelated movie with music.

Two fixed length test trials of both 44s were shown to the infant after the break. In the new action trial the movie in which the cow took the short pathway was shown exactly 4 times. The movie in which the cow took the long pathway took longer than the new action movie. Thus, in the old action trial, the movie was also shown 4 times, but the fourth time the movie stopped earlier. However, the crucial parts of the movie (from the beginning until the cow appears from under the blue oval) are still shown four times.

Note that for the first 6.5 s both the old action movie and the new action movie were exactly the same. Only the period *after* the first 6.5 of the movie was different in the old action and the new action movie. To be sure the looking times at both test trials reflected their reaction to the nature of the stimulus requires the infants to look

in the period of the movie after the first 6.5 s. Therefore, only infants who looked at the important period in at least one of the four movies from each test trial were included in the analysis.

Coding and data analysis

Infants' looking behaviour was recorded at 50 Hz, so we retrieved information every 20 ms about where the infant looked at the screen. We used this data to measure (1) the overall looking time on the first and second test trial was measured (2) the first anticipatory look of the infant to one of the two pathways during the 1.5s period the cow is under the blue oval. To analyse anticipations we calculated a difference score (DS) (cf. Paulus, Hunnius, & Bekkering, in prep; for a similar procedure see Gredebäck, Theuring, Hauf, & Kenward, 2008) we coded the movies as follows: if an infant anticipated to the long pathway this movie was given the value 1, if the infant anticipated to the short pathway the movie was given the value -1. If the infant did not anticipate at all during the time the cow was under the blue oval, the movie was coded with the value 0. The areas in which the gaze of an infant counted as an anticipation had the same surface (139385 px).

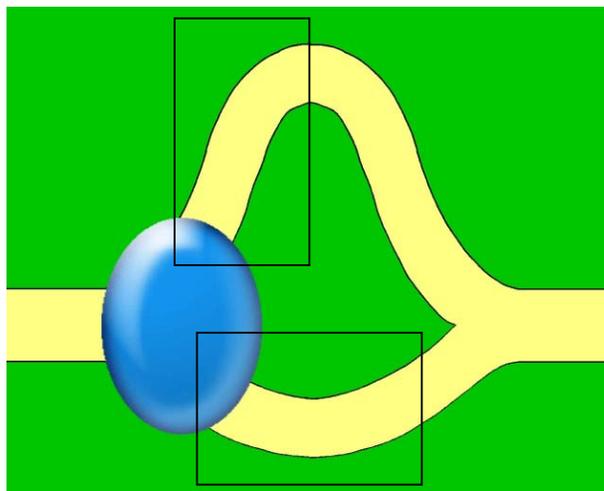


Figure 6: The areas (approximately) in which a gaze of an infant counted as an anticipation. Both areas had the same surface. If the infant gazed at the upper area, the infant is said to be anticipating to the long pathway. If the infant gazed at the lower area, the infant is said to be anticipating to the short pathway.

We analyzed the overall looking time on the two test trials by using a paired samples t-test, pairing the looking times on the Old Action trial with the New Action

trial. The difference score on the first movie of the first trial was analyzed with a one-sample *t*-test. This analysis was the most important, because this was the first movie in which the short pathway was intact and the infants had not saw the cow choosing one of the two pathways yet in this setting. Further we were interested if the infants, who saw the trial in which the cow took the short pathway as their first test trial, changed their anticipation behaviour in the remaining movies of this test trial. This was analyzed with a repeated measures ANOVA with the DS as dependent variable and the within-subjects factor New Action Trial (movie1, movie2, movie3, movie4).

RESULTS

Habituation

Table 2 shows the mean looking times of the first and last three habituation trials and the test trials for the two groups (old action first, new action first). The average number of trials necessary for habituation was 10 (SEM = 0.63).

	Habituation trials						Test trials	
	First three trials			Last three trials			Old action	New action
Old action first (n =9)	18.85	12.43	22.97	12.25	3.47	8.3	26.00 (2.95)	25.82 (3.36)
New action first (n=10)	23.41	24.07	22.11	6.12	2.87	6.03	20.90 (1.63)	23.65 (2.89)

Table 2: The mean looking times in seconds on the habituation trials and the two test trials for both groups (old action first, new action first). The number in de brackets is the standard error of the mean.

A repeated measures ANOVA with dependent variable Looking Time and the between-subjects factor First Action (old action, new action) and within-subjects factor Action Type (old action, new action) revealed no effect ($F(1, 17)=0.672, p>0.1$), so we averaged the looking time of all subjects on the two test-trials. During the test phase the mean looking time to the old action was 24.68s (9.38) and the average looking time new action was 23.32s (SEM = 7.42). A *t*-test for paired samples with Action Type (old action, new action) as independent variable and Looking Time as dependent variable revealed that the action type had no significant effect on mean looking time ($t(18) = .771, p = 0.451$, see Figure 7).

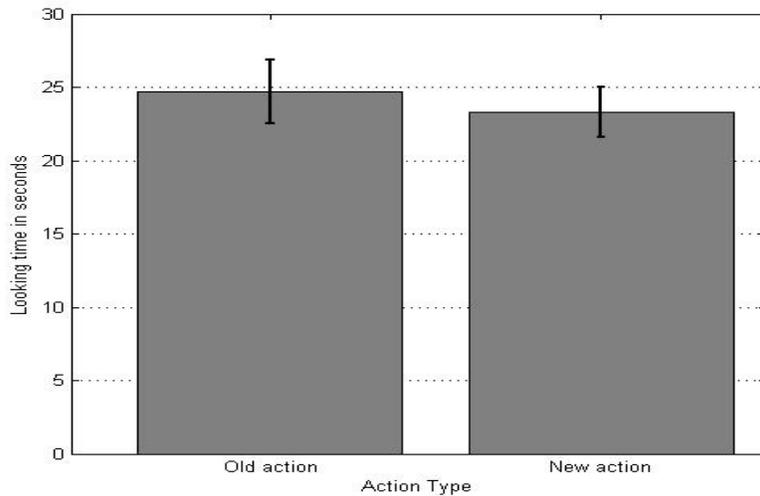


Figure 7: The mean Looking Time in seconds on the test trials (Old action, New action). The black bars stand for the standard error of the mean.

Anticipation

Of all infants 63% anticipated in the first trial and 92% of them anticipated to the long pathway. A one-sample t -test with difference score as variable showed that the difference score was significantly different from zero ($t(18) = 3.750, p = 0.001$, see Figure 8).

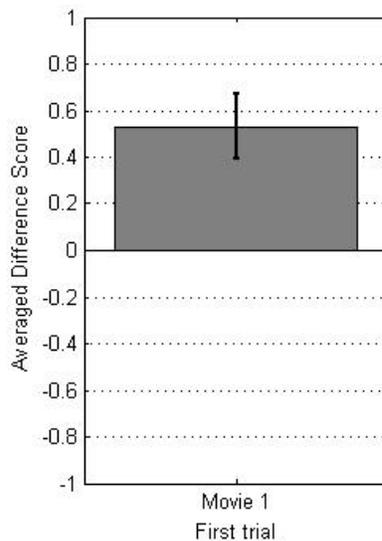


Figure 8: The Averaged Difference Score of all infants on the first movie from the first trial. An Average Difference Score of 1 means all infants anticipated and they all anticipated to the long pathway. An Average Difference Score of -1 means all infants anticipated and they all anticipated to the short pathway. An Average Difference Score of 0 means all infants did not anticipated at al. The black bar in the figure stands for the standard error of the mean.

A repeated-measure ANOVA with dependent variable Difference Score and within-subjects factor New Action Trial (movie1, movie2, movie3, movie4) showed that there was no significant difference in anticipation behaviour in the infants (who saw as their first test trial the trial in which the cow took the short pathway) over the four movies ($F(1, 9) = .778, p = 0.543$ see Figure 9), so we averaged the difference score over all movies.

An one-sample t -test with the Averaged Difference Score as variable was performed and showed that infants anticipated on average significantly more to the long pathway than to the short pathway ($t(9) = 3.000, p = 0.015$).

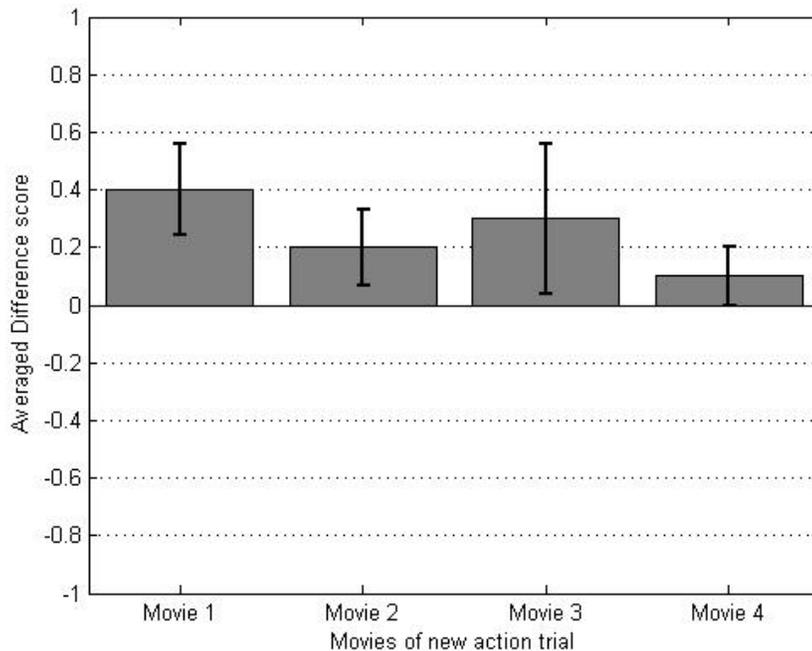


Figure 9: Averaged difference score of the infants who saw the new action as their first test trial. This figure shows the averaged Difference Score at all four movies of this test trial. An Average Difference Score of 1 means all infants anticipated and they all anticipated to the long pathway. An Average Difference Score of -1 means all infants anticipated and they all anticipated to the short pathway. An Average Difference Score of 0 means all infants did not anticipate at all. The black bars in the figure stand for the standard error of the mean.

DISCUSSION

The aim of this experiment was to investigate if action anticipation in infants is based on the rationality or the frequency of an action. Infants were habituated to a movie in which a cow walked on a long pathway towards a sheep as another, shorter

pathway was blocked. In the two test trials both pathways were intact and in one trial the cow took the short pathway, in the other trial the cow took the long pathway. For both test trials the looking time of the infants to these trials were measured as well as the anticipation of the infants to one of the two pathways. The anticipation measurements showed evidence for the frequency hypothesis, the looking time measure, however, yielded no evidence for either the frequency hypothesis or the rationality theory.

Thus, we found strong evidence for the frequency hypothesis by measuring anticipation. Infants anticipated that the cow would take the long pathway, although there was also a shorter pathway to reach the same road. Even after the infants observed the cow taking the short pathway a few times the infants still anticipated to the long pathway. These findings suggest that infants make anticipations about actions of others based on frequency, instead of rationality. This finding suits well with findings from other areas like language acquisition (Gomez, 2002) and visual perception (Fiser, & Aslin, 2002) showing that infants are sensible for statistical probabilities. It can be speculated that their ability to learn from frequency forms the basis of their development of higher order capacities. Further research is needed to address this question.

Currently we are also running the experiment with adults and the preliminary results show that even adults expect the cow to take the long pathway in the first movie of the first test trial. It seems, thus, that even adults base their action anticipation on frequency. However, when the adults were asked about the goal of the cow, the adults could easily give a rational explanation for the fact that the cow took the short pathway. Unlike the infants, when the adults saw the cow take the short pathway once, they expected it to take the short pathway from then on. A possible cause for this difference is that infants base their action anticipations only on frequency, but that adults can switch to rational action anticipations if the on frequency based expectancies are not fulfilled (in this case their expectancy, based on frequency, of the cow taking the long pathway was not fulfilled). Further research is needed to test this hypothesis.

Although we did find evidence for the frequency hypothesis with measuring anticipation, we did not find evidence for either the frequency or the rational theory with measuring the looking time of the infants. According to an interpretation of looking times in habituation experiments, infants look longer to a movie which is in

conflict with their expectations (Spelke, 1985). Since no difference was found in looking time between the two test trials, it appears that infants are evenly surprised by the cow taking the long (irrational) pathway and the cow taking the short (rational) pathway. This finding is in contrast to the findings of Csibra et al. (1999), who found for the same age group that infants looked longer to the irrational action than to the rational action.

One possible explanation for the fact that we did not find evidence for either the frequency hypothesis or the rationality theory in measuring looking time, is that as well as the rationality of an action as the frequency of an action play a role in the action anticipation of an infant. In the experiment of Csibra et al., the irrational action was also the least frequent action, and the rational action was the most frequently seen action. Hence, if the rationality and frequency of an action both play a role in action anticipation it could be expected that infants were surprised by the irrational *and* non-frequent action, and thus looked longer to it. This explanation is also consistent with the finding that infants were not surprised by one of the two actions in our study, because one of the actions was the most frequently seen *but* irrational action, whereas the other action was a rational *but* not frequently observed action. Hence, in our study the high frequency of one action and the rationality of the other action possibly cancel each other out, resulting in similar looking times at both actions.

There is, however, a problem with this explanation. The anticipation measure, on the contrary, revealed evidence in favour of the frequency hypothesis. This suggests an alternative explanation for the finding that infants displayed no difference in looking time for the two experimental conditions: action anticipation in infants *is* based on frequency, but the habituation procedure is not precise enough to measure this. Results based on the habituation procedure turn out to be open for different interpretations (Bogartz et al., 1997; Hunnius, 2007) and it seems that habituation is more suitable to globally investigating perceptual processing in infants. It is possible that the differences in looking time on both test trials in Csibra et al. (1999) are caused by other differences than rationality. For example it is possible that the infants in the experimental condition looked longer to the old action because the action itself is more interesting than the new action. Whereas the new action consists of one continuous action to the other side of the screen, the old action is not continuous and consists of three parts, going forward, going backward, en going forward (jumping) again. In our experiment both actions were continuous. However, in de control

condition of Csibra et al. there was no difference found between the two test trials, although the infants were habituated with the same action as in the experimental condition. This could be explained by the fact that the perceptual difference between the habituation trials and the old action trial was bigger in the experimental condition than in the control condition. In the experimental condition the black block (see figure 1) was more part of the action, than in the control condition. This because the black block in the habituation trials in the experimental condition was in the middle of screen and constantly in sight when the action was performed. In the control condition the black block was at the side of the screen and was out of sight during a big part of the action. Thus, although the old action is more interesting than the new action, when the black block was gone in the test trials, the old action movie seemed newer in relation to the habituation trials for the infants in the experimental condition than for the infants of the control condition. It is possible that these perceptual differences were causing the results of the experiment performed by Csibra et al.

As stated in the introduction, it could be interesting to develop capabilities of a social robot in the same way as we humans develop, by first programming infant-level intelligence and then let this social robot develop either through experience and imitation or by explicit programming different developmental stages (Turing, 1950; Bakker & Kuniyoshi, 1996; Demiris & Meltzoff, 2008). To build a social robot equipped with the capability to anticipate actions of others, in the same way we humans do, we thus need an infant-based action anticipation model. Many existing computational models in Artificial Intelligence concerning action anticipation and goal prediction are based on the assumption that infants expect actions of others to be the most rational action possible to reach a goal (Oztop et al., 2004; Baker et al., 2007; Mao & Gratch, 2004; Youn & Oh, 2008; Verma & Rao, 2006). However, the finding that infants, and possibly even adults, base their anticipation of actions on frequency and not on rationality suggests the need for models based on frequency. Instead of representing beliefs, desires and intentions, a model could be presented with many actions and their goals. The model could extract the similarity and frequency of these actions (for example with the use of neural networks (Patterson, 1998) and use this information to make anticipations about actions of other agents. It may be interesting to test how these models based on frequency perform in relation to models based on rationality. Also, the preliminary results of the adult study imply that it is possible that adults can also base their anticipation of actions on rationality if

their action expectations based on frequency are wrong. For further research it could be interesting to investigate ways in combining a frequency-based action model with a rationality-based action model.

Besides predicting actions of agents in the environment, action models can be used for a more general purpose in the development of social robots, namely learning by imitation. Learning by imitation enables a robot to learn new behaviours by observing (and imitating) other agents in its environment (Bakker & Kuniyoshi, 1996) and the use of this learning method is growing in popularity. Infants are imitating at a very young age (Meltzoff, 1995; Bekkering, Wohlschläger & Gattis, 2000), but they do not imitate every action they witness (Breazeal & Scassellati, 2002). Therefore, one of the problems with imitation in robots is: how does a robot know what to imitate? A good action model could be a solution for the selective imitation problem; if a robot has a good action anticipation model, it would be easier for the robot to separate “noise” actions (actions not important for reaching the goal) from goal-directed actions.

In sum, this study provides evidence for the hypothesis that infants (and possibly even adults) base their action anticipation on frequency. As a result of this finding, it could be interesting to see how action models in Artificial Intelligence based on frequency will perform in comparison with or in cooperation with existing models based on rationality.

The understanding of actions and goals of others helps us humans to communicate and cooperate with each other. It takes several years of observing and imitation to learn, understand and predict actions and their goals. If we want social robots in our environment with which we can cooperate just like humans, I think it is important to keep looking at the way we humans develop in our lifetime and learn from that development.

ACKNOWLEDGEMENTS

First I want to thank my supervisors Iris van Rooij, Markus Paulus and Sabine Hunnius for their help with performing the experiment and advice in writing. I would like to thank Sven Vrins for programming my experiment for the Tobii and for help with problems with using it. I would also like to thank Angela Khadar for general support in the Baby Research Centre in Nijmegen.

REFERENCES

- Baker, C.L., Tenenbaum, J.B. & Saxe, R.R. (2007) Goal inference as inverse planning. *Proceedings of the Twenty-Ninth Annual Conference of the Cognitive Science Society*.
- Bakker, P. & Kuniyoshi, Y. (1996) Robot See, Robot Do: An Overview of Robot Imitation. *AISB96 Workshop on Learning in Robots and Animals*.
- Bekkering, H., Wohlschläger, A., & Gattis, M. (2000). Imitation of gestures in children is goal-directed. *Quarterly Journal of Experimental Psychology*, 53, 153-164.
- Bogartz, R.S., Shinsky, J.L., & Speaker, C.J. (1997). Interpreting infant looking: The event set x event set design. *Developmental Psychology*, 33, 408-422.
- Breazeal, C. & Scassellati, B. (2002). Robots that imitate humans. *Trends in cognitive Sciences*, 16 (11), 481-487.
- Csibra, G. (2003) Teleological and referential understanding of action in infancy. *Philosophical Transactions of the Royal Society, London B*, 358, 447-458.
- Csibra, G., Gergely, G. (2007) 'Obsessed with goals': Functions and mechanisms of teleological interpretation of actions in humans. *Acta Psychologica*, 124, 60-78.
- Csibra, G., Gergely, G., Bíró, S., Koós, O. & Brockbank, M. (1999) Goal attribution without agency cues: the perception of 'pure reason' in infancy. *Cognition*, 72, 237-267.
- Davidson, D. (2001), *Essays on Actions and Events*, Oxford: Clarendon Press, 2nd edition.
- Demiris, Y. & Meltzoff, A. (2008) The robot in the crib: a developmental analysis of imitation skills in infants and robots. *Infant and Child Development*, 17, 43-53.

- Falck-Ytter, T., Gredebäck, G., & von Hofsten, C. (2006) Infants predict other people's action goals. *Nature Neuroscience*, 9 (7), 878-879.
- Fiser, J., & Aslin, R.N. (2002). Statistical learning of new visual feature combinations by infants. *Proceedings of the National Academy of Sciences*, 99, 15822-15826.
- Gergely, G., Bekkering, H., & Király, I. (2002). Rational imitation in preverbal infants. *Nature*, 415, 755.
- Gergely, G. & Csibra, G. (2003) Teleological reasoning in infancy: the naive theory of rational action. *TRENDS in Cognitive Sciences*, 7, 287-292
- Gergely, G., Nádasdy, Z., Csibra, G. & Bíró, S. (1995) Taking the intentional stance at 12 months of age. *Cognition*, 56, 165-193.
- Gomez, R.L. (2002). Variability and detection of invariant structure. *Psychological Science*, 13, 431-436.
- Gredebäck, G., Theuring, C. Hauf, P., & Kenward, B. (2008). The microstructure of infants' gaze as they view adult shifts in overt attention. *Infancy*, 13(5), 553-543.
- Haselager, W.F.G. (1997) *Cognitive Science and Folk Psychology: The Right Frame of mind*. Londen: Sage.
- Hauf, P., & Prinz, W. (2005). The understanding of own and others' actions during infancy: "You-like-me" or "Me-like-you" *Interaction Studies*, 6(3), 429-445.
- Hunnius, S. (2007). The early development of visual attention and its implications for social and cognitive development. *Progress in Brain Research*, 164, 187-209.
- Mao, W. & Gratch, J. (2004) A Utility-Based Approach to Intention Recognition, In *Proceedings of the AAMAS 2004 workshop on Agent*.
- Meltzoff, A. N. (1995). Understanding the intentions of others: Re-enactment of intended acts by 18-month-old children. *Developmental Psychology*, 31, 838-850.
- Oztop, E., Wolpert, D. & Kawato, M. (2004) Mental state inference using visual control parameters. *Cognitive Brain Research*, 22, 129-151
- Paulus, M., Hunnius, S., & Bekkering, H. (in prep.). Gaze following in infancy: How infants relate persons and objects. *Paper in preparation*.
- Paulus, M., Hunnius, S., Vissers, M., & Bekkering, H. (2009) Imitation in 14-month-old infants: Rational or motor resonance? *Paper submitted for publication*.

- Patterson, D.W. (1998) *Artificial Neural Networks: Theory and Applications*, Prentice Hall PTR, USA.
- Pylyshyn, Z. W. (1987). *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*. Norwood, NJ : Ablex Publishing.
- Sabanovic, S., Michalowski, M.P. & Simmons, R. (2006) Robots in the Wild: Observing Human-Robot Social Interaction Outside the Lab. *Proceedings of the International Workshop on Advanced Motion Control, Istanbul Turkey*.
- Sirois, S., & Jackson, I. (2007). Social cognition in infancy: a critical review of research on higher-order abilities. *European Journal of Developmental Psychology*, 4, 46-64.
- Spelke, E.S. (1985). Preferential looking methods as tools for the study of cognition in infancy. In G.Gottlieb & N.Krasnegor (eds.), *Measurement of audition and vision in the first year of post-natal life*, Norwood, NJ: Ablex, 323-363.
- Schwier, C., van Maanen, C., Carpenter, M., & Tomasello, M. (2006)'Rational Imitation in 12-Month-Old Infants', *Infancy*, 10(3), 303-311
- Turing, A.M. (1950). Computing machinery and intelligence. *Mind*, 59, 433-460.
- Verma, D. & Rao, R.P.N. (2006) Goal-Based Imitation as probabilistic Inference over Graphical models. *Advances in Neural Information Processing Systems*.
- Woodward, A.L., Sommerville, J.A., & Guajardo, J.J. (2001). How infants make sense of intentional action. In B.Malle, L. Moses, & D.Baldwin (Eds.), *Intentions and intentionality: Foundations of Social Cognition* (pp. 149-169). Cambridge, MA: MIT Press.
- Youn, SJ. & Oh, KW. (2008) Intention Recognition using Graph Representation, *International Journal of Applied Science, Engineering and Technology*, 4(1).