http://dx.doi.org/10.14236/ewic/HCI2016.28

Children’s lying behaviour in interactions with personified robots

Mariana Serras Pereira, Yoeri Nijs, Suleman Shahid, Marc Swerts
Tilburg center for Cognition and Communication, Tilburg University
Tilburg University, School of Humanities, PO Box 90153, 5000 LE Tilburg, The Netherlands
m.serraspereira@uvt.nl, yoerinijs@gmail.com, s.shahid@uvt.nl, m.g.j.swerts@uvt.nl

This study investigates how young children between 4 - 6 years old interact with personified robots during a lying situation. To achieve this, a temptation resistance paradigm was used, in which children were instructed to not look at a toy (behind their back) while the instructor (a robot dog, a humanoid or a human) left the room. Results revealed that regardless of the type of communication partner, children’s peeking behaviour was similar across the 3 conditions, while there was a tendency of lying more towards the robots. The majority of the children (98%) showed semantic leakage while telling a lie, and most of them (89%) lied and denied their peeking behaviour. Additionally, children generally gave more verbal responses to the robot dog and to the humanoid in comparison with the interaction with the human. Furthermore, the mean pitch of children differed between the robot conditions, i.e. the mean pitch was significantly lower in the robot dog condition in comparison with the humanoid condition. Finally, facial expression analysis showed that children generally appeared happier when they were interacting to the robot dog compared to the humanoid or human.

Children, Human-robot interaction, Lying behaviour, Nonverbal expressions, Verbal expressions

1. INTRODUCTION

Child-robot interaction is an emerging research field, which recently has yielded a significant amount of studies, ranging from supporting teaching and educational activities through robotic applications (Tanaka et al., 2007; Draper & Clayton, 1992) to helping autistic children in training their social skills using social robots (Stanton et al., 2008). In general, robots in this area of research are explicitly designed to build relationships with humans, and to bring an added value to children’s life. Accordingly, social robots are becoming more adaptive, personified, embodied and autonomous (Breazeal, 2003) and have been shown to strongly influence the way children perceive the world (Kahn et al., 2004; Severson & Carlson, 2010; Turkle, 1999).

However, while there may be many benefits of having robots with which one can build close relationships, social robots could also be misused as well. Particularly, social robots can be deceptive towards people, and this artificial deceptiveness can lead to various threats. For instance, people may assume that the confidential information they passed to a robot, remains confidential, but in reality social robots may pass these sensitive and secret details unnoticed to a third party (Westlund & Breazeal, 2015). Therefore, artificial deceptiveness can cause serious security breaches in specific contexts (Coeckelbergh, 2012). Child-robot interaction is one of those areas where this artificial deceptiveness could prove to be particularly vulnerable (Westlund & Breazeal, 2015). The impact that robots might have in children’s life raises questions and concerns about trust and privacy, in particular because according to previous studies children tend to treat robots as friends and companions (Kanda et al., 2004; Kahn et al., 2012).

Since artificial deceptiveness could potentially effect the future use and acceptance of robots by children in everyday life, it is important to understand this phenomenon by conducting more research on how children behave towards robots in deceptive contexts. Because social robots could be deceptive towards children, it is important to investigate to what extent children view different types of robots as trustworthy partners, and how this compares to their interactions with human beings. One intriguing questions in this respect is related to the degree to which the children’s Theory of Mind (ToM), i.e., the ability to attribute mental states to another partner and recognize differences between one’s own and the other’s perspective, varies as a function of the type of communication

Proceedings of British HCI 2016 - Fusion, Bournemouth, UK

partner. Previous research has proven that in order to produce a successful lie, children need, first of all, to understand their own mental state as well as the mental state of the person to who they are lying (first order belief), and also to keep semantic control over the entire lie (second order belief) (Talwar et al., 2007). Additionally, it is known that children leak some verbal and nonverbal cues while telling a lie despite the fact that previous studies showed inconsistent results regarding which cues are the most relevant and reliable for lie detection (Feldman et al., 1979; Lewis et al.,1989; McCarthy & Lee 2009; Talwar & Lee, 2002a ;Vrij et al., 2004). Therefore, it is relevant to explore if children exhibit similar verbal and nonverbal cues while lying to robots, as the ones that they show during human interaction.

In addition, as social robots for children are coming in different forms, with movements, shapes and behaviours that could be artificial or more human-like, it is relevant to explore if children's deceptive behaviour is affected by such variability in the robot's appearance. In sum, the present study aims to gain understanding of children's lying behaviour towards robots, not only to improve child-robot interaction, but also to shed light on human deceptive skills in various contexts.

1.1 Children's Lying behaviour

Children's ability to lie appears as early as 2.5 and 3 years, and tends to improve with age (Lewis et al., 1989; Newton et al. 2000) According to previous studies, learning to lie is an essential step and part of a normative behaviour in children's development (Talwar & Crossman, 2012; Talwar & Crossman, 2011). Around 3 years-old, children have already some conceptual understanding of lying behaviour (Siegal & Peterson,1998; Talwar & Lee, 2002b), probably because in early stages of their life, parents and caregivers taught children the negative moral implications associated with telling a lie (Xu et al., 2010).

Previous work suggests that there are two main types of lie that occur during children's socialization (Xu et al., 2010). The first type consists of so-called antisocial lies. These lies tend to violate moral rules, and are told for self-serving purposes. These are usually the first type of lies that children are able to produce (Lewis et al. 1989; Talwar & Lee, 2002a). The second type consists of prosocial lies (white lies), and these are told with the intention to benefit or help another person, and/or for politeness purposes. Studies about white-lies have shown that children between 3-7 years old are able to produce a white-lie for the benefit of the other (Talwar & Lee, 2002a), and for politeness purposes (Talwar et al., 2007).

In order to produce a consistent lie, children not only need to be able to control their nonverbal behaviour, but also need to avoid what is called semantic leakage. This means that children need to maintain consistency between their initial false statements and follow-up statements in order to produce a reliable lie (Talwar & Lee, 2002a), their nonverbal behaviour should appear as natural comparable to what they show in truthful situations, and not reveal obvious signs of stress, guilt or nervousness. Findings from previous studies regarding such issues of children's nonverbal behaviour during a lie-tell are fairly inconsistent. Some studies have linked more positive nonverbal cues with deception, such as smiles, confident facial expressions and a more positive tone of voice (Feldman et al., 1979; Lewis et al., 1989). Other studies have shown that children have less control over their nonverbal expression while producing a lie (McCarthy & Lee, 2009; Vrij et al., 2004).

Regarding verbal cues on children's deceptive speech, studies brought to light that young children are not the most skilful liars (Talwar & Lee, 2002a;). The findings suggest that between 3-5 years old, children cannot keep their deceptive discourse semantically coherent and consistent with an initial lie. But between 6-7 years old, half of them are able to keep a certain level of semantic leakage control, and consequently diminish the risk of being caught by others.

In sum, the previous findings are quite inconsistent regarding nonverbal and verbal cues that children might leak while telling a lie. Therefore, it is relevant to explore if these possible cues are also shown when children lie and interact with different robots; and if children exhibit similar ToM towards robots, i.e., whether children "beliefs" about a robot's mental state, and how these compare to their beliefs about human communication partners.

1.2. Lie detection methods

Past research on deception in general has shown that the automatic or human detection of lies is a very demanding task, with accuracy levels usually around chancel level (Bond & Depaulo, 2006; Edelstein et al., 2006). A meta review of 125 studies about deception revealed that there is not a single unique verbal, nonverbal or physiological cue related to deception (Vrij, 2004). However, several techniques have been used for lie detection, from human judges that operate as lie detectors to more novel and automated measures, due to the advancement in Social Signal Processing

(SSP) For instance, eye tracking technology has been used to distinguish liar's gaze patterns from truth-tellers (Wang et al., 2010). Automated movement analysis has started to be used for the same purpose (Serras Pereira et al., 2014; Eapen et al., 2010; Duran et al., 2013), as well as physiological data, such as galvanic skin conductance (Van't Veer et al., 2014), and brain activity (Ding et al., 2013; Kozel et al., 2005). However, despite the variety of tools used for detection, there is not yet a clear and systematic way to achieve highly accurate lie detection results. These methods require not only a considerable amount of experimentation, and lie detection methods that rely on only a limited set of features fail to produce good results, since as described above there is a range of possible cues to deception, from verbal to nonverbal signals. Hence, the present study uses multi-method approach for analysing children's interaction and lying behaviour during human-human and human-robot communication.

1.3. Children's beliefs about robots

It is clear that children are quite susceptible to robots, and often tend to treat robots as friends (Kanda et al., 2004; Kahn et al., 2012). Robots can easily gain the trust of young children. It has been argued that, compared to what they do with their puppets, children feel more inclined to share their secrets with humanoid robots (Bethel et al., 2011); as a matter of fact, the use of puppets as a technique to help children sharing their 'secrets' has been shown not to be very efficient or successful (Carter & Mason, 1998; Johnston, 1997). In the study from Bethel et al. (2011), children between 4-6 years old were asked to keep a secret, and later on were prompted either by a humanoid (NAO robot) or a human to tell that secret. Qualitative results indicated that children were as likely to share the secret with the robot as the adult (with a similar amount of prompting effort). Moreover, these children interacted with the humanoid using similar social conventions as observed in their interactions with the adult, such as greeting, turn taking, etc. This finding is interesting in view of the assumption that there might be a disconnection between what children know about the functioning of robots and what they think about robots as entities (Westlund & Breazeal, 2015). For instance, around 4 years children barely attribute any biological property to robots despite the fact that they still attribute some psychological capacities, such as emotions and cognition (Jipson & Gelman, 2007). Children around 5 year-old believed that robots do not have a brain, however children between 7-11 years old assumed that robots have a certain type of brain that is different from the human version (Scaife & van Duuren, 1995). In addition, children who have had experience with robots tend to attribute intelligence features to a robot, instead of aliveness features (Bernstein & Crowley, 2008). Moreover, according to these children this level of intelligence is different and distinct from human or animal intelligence. Similarly, the results also showed that children with (almost) no experience with robots, not only attributed aliveness features, but also emotional and intellectual abilities to robots.

Thus, past work has shown that children have a tendency to attribute some of the human abilities to robots. However, this attitude towards robots may depend on the kind of robot they are interacting with, which may look very humanoid or more artificial in nature (such as robot dogs). In a study that compared children's interactions with a robot dog and a (real) dog, it became clear that children (aged between 7-15 years old) showed closer proximity and more touching with the real dog. However most of the children also treated the robot dog in ways very similar to the interaction with the real dog. Surprisingly, children also attributed mental states (56%), social skills (70%) and moral standing (76%) to the robot dog (Melson et al. 2005). Similarly, in a study that focused on children's reasoning and interactive behaviour towards a robot dog (AIBO robot), 66% of the children accorded mental states, social rapport and moral standings to the robot dog. Furthermore, 50% of the children attributed biological properties and 25% also attributed some animacy properties to the robot dog (Kahn et al., 2006). Likewise, in a different study children have shown to speak similarly to a real dog and to a robot dog. The majority of the children gave commands as frequently to the robot dog as to the real dog. Furthermore, children used body movement and objects such as balls to elicit play with the robot dog (Melson et al., 2009).In other words, we have gained insight into the way children interact with and feel about different types of robots. To explore this further and see to what extent children view robots as trustworthy partners, the current study will look into children's lying behaviour. As discussed above, lying has been argued to be related to children's ToM, and their beliefs about the other's mental states. Little is known about children's deceptive skills towards different types of robots. It might be the case that children's lying behaviour towards robots diverges significantly from how they lie to a human. In addition, their behavioural patterns in deceptive situations may vary as a function of the type of robot: maybe, telling a lie to a humanoid turns out to be similar to a human because humanoids are closer in shape, and children tend to attribute some of the human aspects to these type of robots. On the other hand, telling a lie to a robot dog might be different, as they are viewed more as pets, such as real dogs, and are more fun and relaxing to play with. Therefore, we will explore deceptive behaviour in children's behaviour using a variant of a well-attested paradigm, and compare interactions of these children with humans, humanoids or robot dogs, and see whether these reveal differences in relative amount of lies, and specific verbal, auditory and nonverbal correlates.

2. DATA COLLECTION

2.1 Lie Elicitation Paradigm

Several different paradigms have been used to investigate children's lying behaviour. In particular, some studies have used a modified version of the temptation resistance paradigm (Lewis et al., 1989; Talwar et al., 2007; Talwar & Lee, 2002b; Talwar & Lee, 2002a). In this type of paradigm, children are given the opportunity to spontaneously lie due to the opportunity to commit a transgression. According to a previous study, in which children had the opportunity to peek at a game's answer and lie about it, half of the children between 6 and 11 years old did not resist the temptation and peeked at the answer (Talwar et al., 2007). Additionally, another study has shown that this paradigm works with children around 4 and 5 years, and moreover some of these children lied over their peeking behaviour (Lee, 2013).

Based on this, the present study used a temptation resistance paradigm in a guessing game to elicit deceptive behaviour among children. The guessing game was played in 3 conditions – human condition, humanoid condition and robot dog condition. The reason for having two different personified robots, specifically a humanoid and a robot dog is based on what was above described in the literature studies about personified robots (Bethel et al., 2011; Kahn et al., 2006; Melson et al., 2009). In short, the reason for having a humanoid robot is because it resembles and it is closer in shape to humans. Regarding the robot dog, it is clear that children easily engage with them, and behave towards them in a more playful way as they do with real dogs.

The sequence of events was very similar across the three conditions. However, in the robot conditions the lie elicitation and the guessing game were conducted either by the robot dog or the humanoid (instead of the human experimenter). Below the control condition is explained in full detail.

It was told to each child that they would play a game, in which the child had to try to guess the toy that was placed behind his/her back. To achieve this, the child was seated in a chair and was told to not look at the toy (initially covered with a blanket) that was placed on a table behind his/her back. Before leaving the room, with the excuse that he (the experimenter) forgot a pen, the experimenter removed the blanket, and emphasized once again that while he was away, the child should not look at the toy. Additionally, the experimenter mentioned that after his returning, they would play the guessing game, and the child could get a prize if the toy was guessed correctly.

Subsequently, the experimenter left the room, and was absent for around five minutes. During this time, the child was alone in the room. After this, the experimenter re-entered the room and said that he hoped that the child did not look at the toy. Then, the experimenter initiated the guessing game that consisted of 6 questions related to the object. The questions were as follows: 1. *"Did you peek at the toy?"; 2."Which colour do you think the toy has?"; 3."How does the toy look like?"; 4."The toy is an animal. Which animal is it?"; 5."Can you describe how the animal looks like?"; 6."Which sound do you think the animal makes?.* After asking the questions, the experimenter told the child the game was over and that he could look at the toy. In all cases, the child received a sticker as a reward.

Regarding the robot conditions, the only difference was the robot shape (appearance) - in one condition, it was a Lego Mindstorms EV3 humanoid whereas on the other one it was Lego Mindstorms EV3 robot dog (Figure 1).

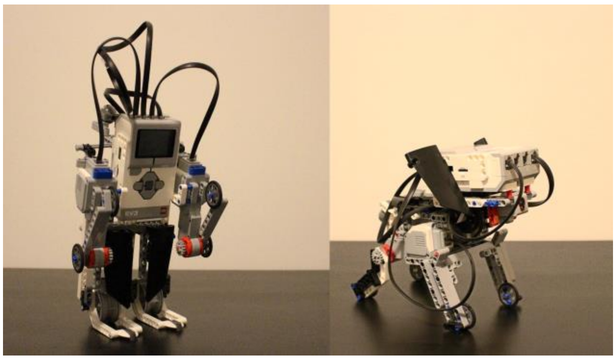

Figure 1. *The Lego Mindstorms EV3 humanoid (left) and the Lego Mindstorms robot dog (right).*

In both robot conditions before the child entered the room, the human assistant asked each child to interact and play the game with the robot. First of all, when the child entered the room, the robot (humanoid or dog) asked the child to sit down on the chair. Further, it told the child that the human assistant would uncover the toy. The robot emphasized that the child should not look at the toy. While the robot was sitting down (and the human assistant was leaving the room), the robot said that its batteries were almost empty, and they needed to be replaced. The human assistant came again and took the robot out of the room. Before leaving, the robot emphasized again that the child should not look at the toy. Like in the control condition, the robot was absent for around five minutes. During this time, the child was alone in the room. When the robot re-entered the room, the game and questions were asked exactly like in the control condition (by the humanoid or robot dog). Additionally, for all robotic statements, a female human voice was used to ensure the robot conditions differ from the human assistant.

2.2 Participants

Eighty-five children from an elementary Dutch school participated (52 boys, 33 girls; mean age = 4.58 years, SD = .60). There were 27 children in the human condition, 28 in the robot dog condition and 30 in the humanoid condition.

2.3 Experimental setup and materials

The children had to sit in front of a table in a room with their back to a second table, where the toy was placed. The toy was a rubber duck (height ≈ 24 centimetres, width≈ 20 centimetres), which was initially and prior to the experiment covered with a blanket. The table in front of the child had a hidden compact camera (Sony NEX-5N), next to another camera (Canon 500D) and two camera bags. During the experiments, the hidden camera (Sony NEX-5N) was making audio and video recordings (Figure 2).

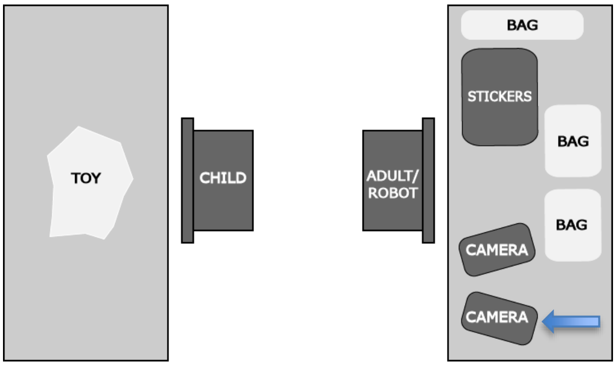

Figure 2: *Experimental setup for the three conditions. The arrow in the figure points to the hidden camera.*

2.4 Procedure

The children were randomly divided in three conditions: the human condition (i.e. the control condition), the robot dog condition and the humanoid condition. The children were told they were going to play a guess game with a hidden toy and they could win a prize if they could guess the correct toy (Figure 3). The experiments were individually paced, and last approximately 10 minutes.

3. RESULTS

The following behaviours - peeking behaviour, verbal behaviour, semantic leakage and nonverbal responses were taken into account for the analysis. The presence or absence of these behaviours was individually coded for each child.

Figure 3: *A child playing the game in the humanoid condition.*

3.1 Peeking behaviour analysis

Regarding peeking behaviour – if children peek at the toy when the experimenter/robots were absent from the room – out of the 85 children that participated, 50 children (59%) were curious during the experiments and showed a clear backward glance. This peeking behaviour – looking at the toy in the absence of the experimenter/robots – was confirmed in the video recordings by the experimenter at a later stage. Furthermore, 63%of the children in the human condition (N=27) peeked at the toy, whereas in the humanoid condition (N = 30) a percentage of 67% looked at the toy, while only 46% in the robot dog condition (N = 28) showed peeking behaviour. To further analyse if there was a significant difference between children's peeking behaviour between the three conditions, a Chi-Square analysis was conducted. Results revealed no statistical significant difference between the conditions.

3.2 Verbal behaviour analysis

The verbal behaviours analysis consisted of two aspects. First of all, the presence of semantic leakage was assessed. An example that clearly illustrates the semantic leakage concept, and that frequently occurred during the experiments was that some children guessed immediately that the toy was a yellow duck. Of course, it was highly unlikely that the children figured it out by themselves, without peeking at the toy while the experimenter/robots were absent from the room. Subsequently, the experimenter asked whether the children peeked at the toy, and most declined and lie about it. The second analysis focused on finding possible differences on children's verbal responses, in particular in exploring possible differences in children's verbal responses between the three conditions, and possible variations in children's mean pitch while interacting with different robots versus a human.

3.2.1 Semantic Leakage

In the present study, from all the children that peeked at the toy when the experimenter/robot was absent (N = 50 in the 3 conditions), the majority (98%) showed semantic leakage. This means that they were inconsistent while producing false statements, showing that they knew information about the object (duck) that they could not guessed unless they looked at the toy. Furthermore, most of them (89%) lied and denied their peeking behaviour, particularly children lied more to the robots (92% in the robot dog condition and 95% on the humanoid condition) than to the human (77%). However, there was no significant difference found in leakage behaviour between the conditions, according to Chi-Square analysis.

3.2.2 Verbal Responses

Regarding the verbal response analysis, the initial focus was whether the children gave or not a verbal answer to the question: *"Did you peek at the toy?"*. Once again, a Chi-square analysis was performed, and revealed a statistical significant difference between the three conditions (x^2(2, N = 85) = 13.29, p < .001). As showed in table 1, it is clear from the results that children gave more verbal responses in both robot conditions (robot-dog: M= .79, SD=. 42 and humanoid M= .77; SD= .43) than in the human condition (M= .37, SD= .49). Additionally, when only focus on the children that peeked at the toy (and also gave a verbal answer to the question), the difference between the three conditions is still significant (x^2(2, N = 50) = 7.40, p < .05) as showed in table 2. Once again, children gave on average more verbal responses towards the robots (robot-dog: M= .85, SD=. 38 and humanoid M= .75; SD=.44 and human: M= .41, SD= .51).

***Table 1**: Chi-square results of the verbal response to the question "Did you peek at the toy?" in the 3 conditions (N=85)*

Condition	N	Mean (SD)	x^2 Stats
Robot dog	*28*	*.79 (.42)*	
Humanoid	*30*	*.77 (.43)*	x^2*(2)= 13.29 ****
Human	*27*	*.37 (.49)*	

****p < .001*

***Table 2:** Chi-square results of the verbal response to the question "Did you peek at the toy?" from the children that peeked at the toy in the 3 conditions (N=50)*

Condition	N	Mean (SD)	x^2 Stats
Robot dog	*13*	*.85 (.38)*	
Humanoid	*20*	*.75 (.44)*	x^2*(2)= 7.40* *****
Human	*17*	*.41 (.51)*	

**p < .05*

These results stimulated further analysis. And therefore, the utterances after the question *"did you peek at the toy?"* were selected for pitch analysis. The means of the pitch values were computed using a Praat script. However because the participants were children, adjustments in the default parameters were performed. The floor was set on 200 hertz while the ceiling was set to 600 hertz, which seem to be the reference values for children pitch analysis (Boersma & Weenink, 2016). The time step used for all the files was 0.01 seconds (i.e. the software computed 100 pitch values per second).

For the current analysis, only the children (N = 55) that gave a verbal response were taken into account. In order to explore possible differences in the mean pitch between the three conditions, a one-way ANOVA was conducted. Interestingly, results revealed a statistical significant difference (F(2, 52) = 7.47, p < .05, η^2 = .223) as depicted on table 3. Tukey post-hoc comparison revealed that the mean pitch in the robot dog (M= 262.88, SD= 37.01) and humanoid (M= 308.00, SD= 46.02) conditions were significantly different, i.e. the mean pitch of the participants in the humanoid condition was significantly higher than the mean pitch of the children in the robot dog condition. No significant differences were found between the robot conditions and the human condition (M= 271. 56, SD= 33.60).

***Table 3**: Anova results from the mean pitch of the verbal answers to the question question "Did you peek at the toy?" in the 3 conditions (N=55)*

Condition	N	Mean (SD)	*F* stats
Robot dog	*22*	*262.88 (37.02)*	
Humanoid	*23*	*308.00 (46.02)*	*F(2, 52) = 7.47**
Human	*10*	*271.56 (33.60)*	

**p < .05*

Furthermore, because there was a significant difference between both robot conditions, it was also analysed if a possible dissimilarity was present when the children's peeking behaviour was taken into account. Once again, a one-way ANOVA was performed, and again there was a strong difference between the robot conditions as shown in table 4 ($F(2, 30) = 7.75$, $p < .005$, $\eta^2 = .341$). Tukey post-hoc comparisons revealed again the trend previously observed – the children from the humanoid condition (M= 314.84, SD= 49.10) had statistically significant higher mean pitch than those in the robot dog condition (M= 252.70, SD=25.17). Once again no statistical significant difference was found between the human condition (M= 275. 37, SD= .39.10) and both robots.

Table 4: *Anova results from the mean pitch of the verbal answers to the question question "Did you peek at the toy?" from the children that peeked at the toy in the 3 conditions (N=33)*

Condition	N	Mean (SD)	*F* stats
Robot dog	*11*	*252.70 (25.17)*	
Humanoid	*15*	*314.85 (49.11)*	*F(2, 30) = 7.75***
Human	*7*	*275.38 (39.11)*	

***p < .005*

3.3 Nonverbal behaviour analysis

This analysis focused on facial cues that children exhibited after the question "*Did you peek at the toy?"*.

3.3.1 Automatic facial expressions analysis

In order to further understand if facial expressions differ across the conditions, an automatic facial analysis was conducted. To achieve this, the Computer Expression Recognition Toolbox (CERT) was used (Littlewort et al., 2011). CERT is designed to automatically detect facial expressions in video sequences. For every frame in a video fragment, CERT calculates the possible presence of the basic emotions – surprise, joy, anger, disgust, fear, sadness and contempt; but it also detects and measures the head pose (yaw, pitch, and roll), and the presence of 30 action units (AU's) from the Facial Action Coding System (Ekman, 1976).

For the present analysis, the same video clips that follow the question "*Did you peek at the toy?"* were used. For CERT analysis, the prototypical emotions, and specifically, four action units were selected. The AU's used were - cheek raise (AU 6), chin raise (AU 17), lip tightening (AU 23) and lip pressor (AU 24). The reason for this was because literature has shown that these AU's are the most reliable indicators of deception in adults (DePaulo et al., 2003).

For every frame in each video clip, CERT calculated the possible presence of the prototypic emotions and action units. After that, for each clip, the mean probability regarding the presence of the basic emotions and AU's was computed. The reason for this was related with an attempt to improve any possible CERT flaws during the data extraction due to, for instance, quick movements or possible blurriness in the clips.

According to table 5, a one-way ANOVA showed that there was a significant effect in AU 6 - cheek raise ($F(2, 49) = 3.91$, $p < .05$, $\eta^2 = .137$), when consider all the children that gave a nonverbal reaction independent of having peeked or not at the toy. A Tukey post-hoc comparison showed that AU 6 (cheek raise) in the robot dog condition (M= .25, SD= .28) was significantly more present than in the human condition (M= .04, SD= .19). However, no significant difference was found between the humanoid (M= .11, SD= .18) and the robot dog, or with the human condition. Additionally, all other action units did not showed statistical significant differences between the conditions.

Table 5: *Anova results for the AU6 - cheek raise - from children that gave a nonverbal answer to the question "Did you peek at the toy?" in the 3 conditions (N=52)*

Condition	N	Mean (SD)	*F* Stats
Robot dog	*17*	*0.25 (0.28)*	
Humanoid	*16*	*0.11 (0.18)*	*F(2, 49) = 3.91**
Human	*19*	*0.04 (0.19)*	

**p < .05*

In addition, a statistical analysis was performed in order to explore possible differences in emotions across the 3 conditions. Interestingly, as shown in table 6, only joy appeared to have a significant effect according to a one-way ANOVA ($F(2, 49) = 4.80$, $p < .025$, $\eta^2 = .171$), while no distinction was made between the children that peeked at the toy and the ones that did not. A Tukey post-hoc comparison revealed a significant difference between the robot dog (M= .03, SD= .04) condition and the human condition (M=.00, SD= .00). However no differences were found between the humanoid (M= .01, SD= .02) and the other two conditions.

Table 6*: Anova results of joy from children that gave a nonverbal answer to the question "Did you peek at the toy?" in the 3 conditions (N=52)*

Condition	N	Mean	*F* Stats
Robot dog	*17*	*0.03 (0.04)*	
Humanoid	*16*	*0.01 (0.02)*	*F(2, 49) = 4.80**
Human	*19*	*0.00 (0.00)*	

**p < .05*

Lastly, the presence of action units and emotions was further investigated within the children that only responded nonverbally to the question about whether they had peeked at the object. As shown in table 7, a one-way ANOVA revealed that the presence of joy differ between the robot dog and the human condition: F(2, 25) = 5.82, p < .05, η^2 = .313. Tukey pairwise comparison showed that in the robot dog condition (M=. .04, SD= .05) children seemed to express more joy compared to the human condition (M= .00, SD= .00).

Table 7: *Anova results of joy from children that only gave a nonverbal answer to the question " Did you peek at the toy?" in the 3 conditions (N = 28).*

Condition	N	Mean (SD)	*F* Stats
Robot dog	*6*	*0.04 (0.05)*	
Humanoid	*6*	*0.01 (0.01)*	*F(2, 25) = 5.82**
Human	*16*	*0.00 (0.00)*	

**p < .05*

4. DISCUSSION

The main goal of the current study was to compare children's lying behaviour in interactions with different types of robots, and with human partners. To achieve this, a temptation resistance paradigm was used, which was inspired by previous work (Lewis et al., 1989; Talwar et al., 2007; Talwar & Lee, 2002b; Talwar & Lee, 2002a). In the present study, children were given the opportunity to peek at a toy (although they were told not to peek), and to lie about their behaviour in order to win a prize. Of the 85 children with a valid response, 50 (58.8%) showed a backward glance. This result is very similar to a previous study in which children showed a similar percentage of peeking behaviour (Talwar et al., 2007). Moreover, this study showed that there is no difference when children interact with humans or with different types of personified robots regarding peeking behaviour, i.e. children in all conditions peeked as frequently to the toy. However, despite the similarity in peeking behaviour, we observed that children lied more to the robots then to human experimenter. Because lying has a moral (negative) valence attached to it (Talwar & Lee, 2002a), children might have considered that lying to robots was less harmful than lying to humans. In addition, children might have assumed that robots could not detect the lies so easily as humans. Furthermore, interacting with robots was probably more playful than interacting with a human, which could have taken away the negative valence of lying, and thus diminishing the threshold for children to lie.

In addition, the experiments in this study confirmed the presence of semantic leakage during a lie-tell. Semantic leakage means that during a lie-tell, children find it hard to keep the information of the initial lie consistent with follow-up statements. From all the children that peeked at the toy (N = 50), a majority of 98% was inconsistent in reproducing a false statement after lying about their peeking behaviour, independent of the condition. And most of them (89%) lied about their peeking behaviour. These results go in line with a previous study, in which children between 3-5 years old showed a poor control of semantic leakage (Talwar and Lee, 2002a). Furthermore, this lack of semantic control provides evidence that in order to successfully lie, children need to have their first order and second order beliefs in ToM fully developed. The semantic leakage found in the present study shows that children between 4-6 years old do not have ToM completely developed, which is also supported by earlier findings (Talwar et al., 2007).

Turning to the results of the verbal analysis, we found that children of 4 -6 years old gave more verbal responses to robots in comparison with the human condition. A possible explanation for this might be that the robots were rather static and gave less interactive cues (e.g. facial expressions and body expressions), and that children therefore tried to overcompensate this lack of feedback by their responses in order to convince the robots (dog or humanoid) of their desired behaviour.

This study also showed that the children's mean pitch differed between the robot conditions. This result goes partly in line with previous findings that suggested that pitch can change during a lie-tell (Streeter et al., 1977). One unanticipated finding was that, in the humanoid condition, the mean pitch was significantly higher in comparison with the robot dog condition. A possible explanation for this might be that children did not feel a strong need to convince the robot dog, because it is not human, and does not resemble any human shape. And there lower pitch could also be interpreted as a sign that children were more relaxed during interactions with the robot dog.

According to this study, children between the ages of 4 - 6 years showed more joy when interacting with robots, specifically when interacting with the robot dog (in line with the pitch results). It can be argued that children found the robot dog playful and were happy while interacting with it. Consequently, the seriousness of the experiment might have been taken away because of the particular shape of the robot, i.e. a dog. Furthermore, these results seem to be consistent with previous findings that found that children that showed an interest in a Lego robot also enjoyed interacting with it (Cook et. al, 2011). Moreover, the results indicate that children tend to attribute social features to a robot dog, which is in line with a previous finding (Melson et al., 2009). Likewise, the new finding about the variation in AU 6 – cheek raise, supports the findings about joy while interacting with robots. Au 6 is one of the AUs that signals happiness/joy. This finding is also in line with previous studies about children's lying behaviour, in which children showed more positive nonverbal cues during a lie tell, such as smiles and a more positive attitude (Feldman et al., 1979; Lewis et al., 1989).

Finally, one limitation of this study is that the robots used were LEGO EV3, and children may have seen the robots as toys (because it is made of Lego); and maybe not as fully autonomous entities because there were some flaws in terms of full interactivity. For instance, they were not build up for rich conversations, and they were not able to (re-)enter the room autonomously (see methodology).

5. CONCLUSION

The present study led to a series of new findings regarding the way children interact with robots, how this compares with humans, to what extent the robot type matters, and how children attribute specific mental states to their artificial and human partners. In particular, we have explored deceptive interactions in various interaction types, which revealed differences in correlates of trust and behavioural patterns. More specifically, the present outcomes of our study contribute to the understanding of child-robot interaction, and to the comprehension of children's deceptive skills towards robots. Furthermore, the findings have significant implications for the understanding of how robots can be used for lie elicitation and lie detection, specifically with children.

6. FUTURE RESEARCH

In future studies it would seem useful to explore a wider range of audio cues and their validity for lie detection, since this study demonstrated that children tended to show an abundance of verbal cues, especially when talking to robots. Finally, it would seem a nice idea to explore whether children behave differently towards other types of robot as well (such as NAO and ICat), given our result that children's beliefs about robots and how they deceive to them may vary as a function of the shape, appearance and human-like features of the robot partner. Finally, the findings about children's ToM towards the robots can also be a valuable insight when designing robots that can be involved in children's daily tasks, such as the ones involved in persuasive games and learning tasks.

6. REFERENCES

Bethel, C.L., Stevenson, M.R. & Scassellati, B., 2011. Secret-sharing: Interactions between a child, robot, and adult. *Conference Proceedings - IEEE International Conference on Systems, Man and Cybernetics*, pp.2489–2494.

Bond, C.F. & Depaulo, B.M., 2006. Accuracy of Deception Judgements. *Personality and Social Psychology Review*, 10(3), pp.214–234.

Ding, X.P. et al., 2013. Neural correlates of spontaneous deception: A functional near-infrared spectroscopy (fNIRS)study. *Neuropsychologia*, 51(4), pp.704–712.

Draper, T.W. & Clayton, W.W., 1992. Using a personal robot to teach young children. *The Journal of genetic psychology*, 153(3), pp.269–73.

Duran, N.D. et al., 2013. Exploring the movement dynamics of deception. *Frontiers in psychology*, 4, p.140.

Eapen, N.M. et al., 2010. The Bodily Movements of Liars. In *33rd Annual conference of the Cognitive Science Society*.

Edelstein, R. S., Luten, T. L., Ekman, P., & Goodman, G.S., 2006. Detecting lies in children and adults. Law and Human Behavior, 30, 1, 1-10. *Law and Human Behavior*, 30(1), pp.1–10.

Feldman, R.S., Jenkins, L. & Popoola, O., 1979. Detection of deception in adults and children via facial expression. *Child Development*, 50 , pp.350–355.

Jipson, J.L. & Gelman, S.A., 2007. Robots and rodents: Children's inferences about living and nonliving kinds. *Child Development*, 78(6), pp.1675–1688.

Kahn, Jr., P.H. et al., 2006. Robotic pets in the lives of preschool children. *Interaction Studies*, 7, pp.405–436.

Kozel, F.A. et al., 2005. Detecting deception using functional magnetic resonance imaging. *Biological psychiatry*, 58(8), pp.605–13.

Lewis, M., Stanger, C. & Sullivan, M.W., 1989. Deception in 3-year-olds. *Developmental Psychology*, 25(3), pp.439–443.

Littlewort, G. et al., 2011. The computer expression recognition toolbox (CERT). *Face and Gesture 2011*, pp.298–305.

McCarthy, A. & Lee, K., 2009. Children's knowledge of deceptive gaze cues and its relation to their actual lying behavior. *Journal of experimental child psychology*, 103(2), pp.117–34.

Melson, G.F. et al., 2009. Implications for the human - Animal bond and for human relationships with personified technologies. *Journal of Social Issues*, 65(3), pp.545–567.

Melson, G.F. et al., 2005. Robots as dogs? Children's interactions with the robotic dog AIBO and a live australian shepherd. *Computer-Human Interaction (CHI) Conference 2005*, pp.1649–1652.

Newton, P., Reddy, V. & Bull, R., 2000. Children's everyday deception and performance on false-belief tasks. *British Journal of Developmental Psychology*, 18(2), pp.297–317.

Scaife, M. & van Duuren, M., 1995. Do computers have brains? What children believe about intelligent artifacts. *British Journal of Developmental Psychology*, 13(4), pp.367–377.

Serras Pereira, M. et al., 2014. Are You Lying to Me ? Exploring Children ' s Nonverbal Cues to Deception. In *36th Annual conference of the Cognitive Science Society*. pp. 2901–2906.

Severson, R.L. & Carlson, S.M., 2010. Behaving as or behaving as if? Children's conceptions of personified robots and the emergence of a new ontological category. *Neural Networks*, 23(8-9), pp.1099–1103.

Siegal, M. & Peterson, C.C., 1998. Preschoolers' understanding of lies and innocent and negligent mistakes. *Developmental Psychology*, 34(2), pp.332–341.

Stanton, C.M. et al., 2008. Robotic animals might aid in the social development of children with autism. *HRI '08: Proceedings of the 3rd ACM/IEEE international conference on Human robot interaction*, pp.271–278.

Talwar, V. & Crossman, A., 2011. *From little white liars to filthy liars: The evolution of honesty and deception in young children* 1st ed., Elsevier Inc.

Talwar, V. & Crossman, A.M., 2012. Children's lies and their detection: Implications for child witness testimony. *Developmental Review*, 32(4), pp.337–359.

Talwar, V., Gordon, H.M. & Kang, L., 2007. Lying in the Elementary School Years: Verbal Deception and Its Relation to Second-Order Belief Understanding Victoria. , 43(3), pp.804–810.

Talwar, V. & Lee, K., 2002a. Development of lying to conceal a transgression: Children's control of expressive behaviour during verbal deception. *International Journal of Behavioral Development*, 26(5), pp.436–444.

Talwar, V. & Lee, K., 2002b. Emergence of White-Lie Telling in Children Between 3 and 7 Years of Age. *Merrill-Palmer Quarterly*, 48(2), pp.160–181.

Tanaka, F., Cicourel, A. & Movellan, J.R., 2007. Socialization between toddlers and robots at an early childhood education center. *Proceedings of the National Academy of Sciences of the United States of America*, 104(46), pp.17954–17958.

Van't Veer, A.E. et al., 2014. Registered report: Measuring unconscious deception detection by skin temperature. *Frontiers in Psychology*, 5(May), pp.1–9.

Vrij, A. et al., 2004. Detecting Deceit Via Analyses of Verbal and Nonverbal Behavior in Children and Adults. *Human Communication Research*, 30(1), pp.8–41.

Wang, J.T., Spezio, M. & Camerer, C.F., 2010. Pinocchio ' s Pupil : Using Eyetracking and Pupil Dilation To Understand Truth-telling and Deception in Games. *The American Economic Review*, 3, pp.984–1007.

Westlund, J., & Breazeal, C. (2015). Deception, Secrets, Children, and Robots: What's Acceptable? Retrieved from http://www.openroboethics.org/hri15/wp-content/uploads/2015/02/Mf-Westlund.pdf

Westlund, J. and Breazeal, C. 2015. The interplay of robot language level with children's language learning during storytelling. Proceedings of the 2015 ACM/IEEE International Conference on Human-Robot Interaction.