2018-08-22

Children conform, adults resist: A robot group induced peer pressure on normative social conformity

Vollmer, A-L

http://hdl.handle.net/10026.1/12259

10.1126/scirobotics.aat7111
Science Robotics
American Association for the Advancement of Science

This is the final draft post-refereeing of Vollmer, A. L., Read, R., Trippas, D., & Belpaeme, T. (2018). Children conform, adults resist: A robot group induced peer pressure on normative social conformity. Science Robotics, 3(21), eaat7111.

For a copy of the pdf of the published paper please visit `http://robotics.sciencemag.org/content/3/21/eaat7111`, or email a request to avollmer@techfak.uni-bielefeld.de or tony.belpaeme@ugent.be.

Children conform, adults resist: robot group induced peer pressure on normative social conformity

Anna-Lisa Vollmer,[1*] Robin Read,[2] Dries Trippas,[3] Tony Belpaeme[2,4*]

[1]Bielefeld University, Cluster of Excellence Cognitive Interaction Technology, 33619 Bielefeld, Germany
[2]Plymouth University, Centre for Robotics and Neural Systems, Plymouth, PL4 8AA, United Kingdom
[3]Max Planck Institute for Human Development, Center for Adaptive Rationality, 14195 Berlin, Germany
[4]Ghent University, IDLab – imec, B-9052 Ghent, Belgium

*To whom correspondence should be addressed; E-mail: avollmer@techfak.uni-bielefeld.de.

People are known to change their behavior and decisions in order to conform to others, even for obviously incorrect facts. Due to recent developments in artificial intelligence and robotics, robots increasingly are found in human environments and there they form a novel social presence. It is as yet unclear if and to what extent these social robots are able to exert similar peer pressure. This study uses the Asch paradigm which shows how participants conform to others while performing a visual judgment task. We first replicate the finding that adults are influenced by their peers, but show that they resist social pressure from a group of small humanoid robots. Next, we repeat the study with 7 to 9-year old children and show that children do conform to the robots. This raises opportunities as well as concerns for the use of social robots with young and vulnerable cross-sections of society; while conforming can be beneficial, the potential for misuse and the potential impact of erroneous performance

cannot be ignored.

One-sentence summary

Children show increased yielding to social pressure exerted by a group of robots, adults however resist being influenced by our robots.

Introduction

Social robots represent a new frontier in the personal robotics industry. These robots are designed to autonomously interact with people across a variety of different application domains in natural and intuitive ways, using the same repertoire of social signals used by humans (*1–3*). Current applications include robotic tour guides in museums (*4*), therapeutic aids in care homes (*5*) and early years childcare (*6, 7*), and teaching aids in primary school classrooms (*2, 8, 9*), with future applications forecast to be far broader (*10*). With these future applications, robots will share the same physical and social space as users, which raises questions regarding safety, and given the social nature of the robots, the psychosocial impact.

It has been shown that people, particularly the younger age groups, easily form strong bonds with social robots, so much so that it can cause distress when a robot is mistreated or misbehaves (*6, 11*), even when they are crude approximations to real living organisms (*12*). Conversely, interaction with social robots has also been found to elicit and reinforce healthy social behaviors in children with autism spectrum disorder (*13–15*) as well as promote and augment social behavior and bonding between group members in care homes (*5*). An open question is whether these social bonds offer robots other affordances such as the ability to exert social influence (*16*), and whether people yield to these.

The *computers as social actors (CASA)* hypothesis (*17–19*) states that people naturally and unconsciously treat computers and other forms of media in a manner that is fundamentally

social, attributing human-like qualities to technology. It has had a notable impact in the fields of Human-Computer Interaction (HCI) and Human-Robot Interaction (HRI). Assuming that the CASA hypothesis holds true, it predicts that people, regardless of their age, are sensitive to (and submit to) social influences exerted by social robots and (crucially) that this is automatic and involuntary (*18*). We tested this prediction by replicating the influential paradigm to study normative social conformity devised by Solomon Asch (*20–22*).

Computers as social actors

Reeves and Nass concluded from a number of social psychology experiments that "individuals' interactions with computers, television, and new media are fundamentally social and natural, just like interactions in real life." (*17, p. 5*). The CASA hypothesis is part of the Media Equation hypothesis (*17*), an overarching theory which additionally implies that people process experiences mediated by technology in the same way as they process unmediated experiences. Describing an unconscious and automatic response, the CASA hypothesis seems to apply to everyone regardless of expertise.

The studies conducted by Reeves and Nass show that people treat technology like people, using the same social rules, expectations, beliefs and behaviors towards technology as they would with other people, according them social behaviors (e.g., politeness, reciprocity), attributing human characteristics to them (e.g., gender), reacting to them as they would to human interaction partners, and so on (*18, 19*). Nass and colleagues found that when a computer asks a user to evaluate itself, the user will give more positive feedback than when the user does the evaluation on a different computer (*23*). They also found that people showed gender stereotypes toward computers with male and female voice (*24*). Rules of attraction seem to hold as well. Users were shown to like electronic partners better when they have the same personality as the user (*17*).

Peer-driven normative conformity and the Asch paradigm

Conformity describes the behavior of an individual who is complying with group norms. In the field of social psychology, two main varieties of conformity are considered: informational social conformity and normative social conformity. The former depicts the influence of others' responses as a source of information on one's own judgment when a task is ambiguous and the correct answer not straightforward. The latter describes an influence of others on judgments in a task with unambiguous stimuli where the correct answers are clear. Participants are lead to give incorrect responses complying publicly with an erroneous majority in order to be accepted.

The well-established and most influential paradigm to study normative social influence was devised by Solomon Asch in 1951 (*20*). In his classic conformity experiments, individual participants were unknowingly grouped with multiple confederates and instructed to judge the length of a target line compared to three comparison lines, only one of which has the same length as the target line (Fig. 1D). For each such comparison, all the participants verbally reported one after the other which comparison line they perceived to match the target line, with the subject verbalizing their answer before the last of the confederates. On two-thirds of the trials the confederates unanimously announced an incorrect judgment (critical trials, $n = 12$) while providing the correct response on the remaining trials (neutral trials, $n = 6$). The participants followed the group response, complying publicly and submitting to group pressure in 32% of trials (in 68% of critical trials they responded correctly; one fourth of the participants were completely independent and resisted the group pressure in all critical trials) (*20*). Asch conducted his first experiment with male college students and a majority group of varying size.

Many replications and alterations of this standard experiment have been conducted to identify factors that influence conformity. Size, immediacy, unanimity, and personal importance of the group, the ambiguity and public announcement of responses, gender, and age are among these factors. Whereas conformity seems to increase with a larger majority, it changes only lit-

tle from group sizes of four (*20, 22*). Majority groups that are personally more important to the participant (e.g. peers, in-group vs. out-group members) (*25, 26*) exert a greater social pressure. If there is only one dissenter in the majority group who announces the correct or even only a different answer from the group, conformity decreases drastically (*21*). It increases as the correct judgment becomes more ambiguous (e.g., by making the line lengths more similar) (*22*). Participants that write down their judgments privately tend to resist group pressure (*22*). Female participants were found to endorse the group response slightly more often than male participants (*27, 28*). Age has been reported to reduce susceptibility to social influence (*29, 30*), although findings seem to be conflicting (*31, 32*).

Results

We have tested whether adults (Experiment 1) and children (Experiment 2) exhibit normative social conformity (*16*) when conducting a visual discrimination task in the presence of three humanoid robots (Fig. 1, A–C). We replicated the Asch paradigm to study normative social conformity. The original group setup formed the basis of our experimental condition. As a control condition, participants were asked to perform the same task while alone. Decreased accuracy on the critical trials in the experimental condition compared to the control condition is evidence for social conformity.

FIGURE 1 ABOUT HERE

Adults

In Experiment 1 we tested the hypothesis that humanoid robots exert normative social pressure on adults. Participants ($N = 60$, 34 female, age: range = 18 – 69 years, $M = 30.9$ years, $SD = 14.2$) were randomly assigned to one of three conditions: a control condition ($n = 20$), a 'human peer' condition ($n = 20$) with three human confederates, and a 'robot peer' condition

($n = 20$) in which three humanoid robots replaced the human confederates.

In all conditions, participants, including the confederates in the human-peer condition and robots in the robot-peer condition, were asked to verbally report which line matched the reference line. The experimenter decided on the response order.

On each trial we measured whether the real participant's verbal response was correct. The experiment was a 3 (condition: control vs. human peer vs. robot peer, between subjects) $\times$ 2 (trial type: critical vs. neutral, within subjects) mixed design. If people are influenced by social peers, line judgment accuracy in the critical (but not the neutral) trials should be lower for the peer conditions compared to the control condition.

Analysis of Logistic Regression model

There was a significant main effect of condition ($\chi^2(2) = 11.8$, $P = .003$), suggesting that peers influenced line judgment accuracy. The condition main effect was qualified by an interaction with trial type, $\chi^2(2) = 11.9$, $P = .003$, indicating that the effect of peers differed for the critical and neutral trials. Follow-up logistic regressions for the critical and neutral trials separately indicated that the presence of human peers significantly reduced judgment accuracy on the critical trials, log-odds = -1.64, $SE = 0.30$, $z = -5.46$, $P < .00001$. No such effect was present for the robot peers, log-odds = 0.26, $SE = 0.37$, $z = 0.71$, $P = .48$. For the neutral trials, there were no significant differences between the conditions: control-human, log-odds = -0.30, $SE = 0.31$, $z = -0.97$, $P = .33$; control-robot, log-odds = -0.03, $SE = 0.32$, $z = -0.09$, $P = .93$. No other effects approached significance, $P > .91$. Accuracy patterns can be found in Fig. 2A. We also found that in the human-peer condition, 83% of the incorrect responses were the same as the confederate response ($\chi^2(1) = 15.114$, $P < .001$), indicating that participants were indeed conforming to the group response (Fig. 3).

This replicates the classical findings of Asch (*20–22*) and confirms recent studies (*33*). Im-

portantly, the drop in judgment accuracy with human peers was present exclusively for the critical trials, suggesting that the performance drop is not due to domain general anxiety driven by the presence of peers.

FIGURE 2 ABOUT HERE

Children

Adults do not appear to normatively conform to the humanoid robots used in the study, providing a challenge to the CASA hypothesis. However, since children are known to be more susceptible to social influence (*29, 30, 34, 35*), we evaluate this finding with young children in Experiment 2. Given the practical challenges of experiments using the original Asch paradigm involving child confederates, we focused exclusively on the influence of humanoid robot peers (cf. Section Outlook).

Participants ($N = 43$, 22 female, age: range = 7 – 9 years, $M = 8.5$ years, $SD = 0.5$) were randomly assigned to either the control ($n = 21$) or robot-peer ($n = 22$) condition. The methods and materials were identical to those from Experiment 1, with the exception that children were tested at school, rather than in a university lab.

We measured children's performance at the task when alone and when in the presence of robots using a 2 (condition: control vs. robot peer, between subjects) $\times$ 2 (trial type: critical vs. neutral, within subjects) experimental setup.

Analysis of Logistic Regression model

The analysis revealed that children are significantly influenced by the presence of robot peers (significant interaction between the two factors, condition and trial type, $\chi^2(1) = 11.1$, $P = .0009$). An analysis of the critical and neutral trials separately indicated that line judgment accuracy was lower in the robot-peer condition than in the control condition for critical trials

(log-odds = -0.37, $SE = 0.12$, $z = -3.17$, $P = .002$) but not the neutral trials (log-odds = 0.21, $SE = 0.15$, $z = 1.4$, $P = .16$). No other effects approached significance (all $P's > .30$). Accuracy patterns can be found in Fig. 2B and Table S1. We also found that in the robot-peer condition, 74% of the incorrect responses during the critical trials were identical to the responses provided by the robots ($\chi^2(1) = 14.785$, $P < .001$), again suggesting that conformity to the majority was taking place (Fig. 3).

FIGURE 3 ABOUT HERE

Discussion

It appears that adults in our study do not conform to the group of robots, confirming recent studies (*33*). Brandstetter et al. used four Nao humanoid robots to investigate informational and normative social influence in adults. The robots in their experiment were individualized with outfits and played pre-recorded human voices in order to focus on the appearance of the robots. Their setup also differed to ours in the length, presentation and number of stimuli. In 33 trials, Brandstetter et al. projected the lines of length up to 110*cm* onto a projection area and found that adult participants were influenced by their peers but not by the robots (neither with ambiguous nor unambiguous stimuli).

Children in our study on the other hand seem to conform to the robots. An alternative explanation for the findings is that children were not influenced or conforming, but rather that the relative novelty of the situation led to an overall decrease in judgment accuracy. This criticism holds no ground, as there was no accuracy decrease for the neutral trials. In fact, if anything, children performed slightly better for such trials (although this finding was not statistically significant), again indicating that they followed the suggestions made by the robots.

There is also the possibility that children were conforming to the robots' responses due to the authority invested in the robots by the adult experimenter. Even so, this still suggests that

the robots exert peer pressure and does not invalidate the observations and conclusions. Robots are likely to be owned by someone, people or organizations, and might as such be proxies for indirect social peer pressure.

The results of these experiments have both theoretical and practical implications. From a theoretical perspective, our results counter the notion that is central to the CASA hypothesis – that all people instinctively and automatically treat computer-based media as social (*17, 18*). While in certain tasks, adults do attribute human-like qualities to machines (*17*), they are capable of inhibiting the effects of normative influence, something which is not observed for human peers. We see this as a refinement of the CASA hypothesis, which impacts on the design of human-machine interaction in general.

Recent studies of online social networks have revealed that user behavior and decision making can be altered and manipulated through the selection of presented information (*36, 37*). Social robots are yet another social medium through which information may be transferred and communicated, and if trusted they can assert informational influence (*38*). The fact that robots have the power to induce conformity, even just in children, is relevant here and we believe our results are both timely and critical. In this light, care must be taken when designing the applications and artificial intelligence of these physically embodied machines, particularly as little is known about the long-term impact that exposure to social robots can have on the development of children and vulnerable sections of society (*39*). More specifically, problems could originate not only from intentional programming of malicious behavior (e.g. robots that have been designed to deceive) but also from the unintentional presence of biases in artificial systems (*40*) or the misinterpretation of autonomously gathered data by a learning system itself. For example, if robots recommend products, services or preferences, will compliance and thus convergence be higher than with more traditional advertising methods?

From a practical perspective, given that children do conform to erroneous suggestions made

by social robots, concerns are raised when using social robots with young people; while conforming can be beneficial (*41, 42*) (for example in health care or education), the potential for misuse or erroneous use cannot be ignored. This is a salient issue as there is a growing interest from the private/industrial sector in robots that interact with the general public and in particular with children. As this industrial market grows, so do the number of children potentially exposed to the issues outlined here.

A future in which autonomous social robots are used as aids for education professionals or child therapists is not distant. In these applications the robot is in a position in which the information provided can significantly impact the individuals they interact with. A discussion is required on whether protective measures, such as a regulatory framework, should be in place that minimize the risk to children during social child-robot interaction and what form they might take as not to adversely impact the promising development of the field.

Outlook

We conducted our experiment with children aged between seven and nine years. To create a more complete picture of conformity to robots, studies with different age groups, including older ages, need to be conducted such that the age ranges in which children and adults conform to robots can be determined.

Conducting the Asch experiment with children is difficult, as all but one of the children need to be confederates and convincingly act as fellow participants. Most studies on conformity with children have thus used a different paradigm to study conformity or used special optical setups giving the participant a different visual experience without the participant realizing (*35, 43*). A human-peer condition with children would have allowed a direct comparison between the results in the human peer condition and in the robot peer condition. The lack thereof, however, is a limitation of the current study.

A review of 133 Asch replication studies shows that conformity in adults has decreased since the 1950s (*28*). In addition, there is a correlation with a society's individualistic or collectivist nature. Compliance on the Asch paradigm is higher in societies with high collectivism, and it would be interesting to see if children and adults in collectivist cultures are more likely to yield to robots than individuals from individualistic cultures.

The sample sizes in our study are limited. Although sample sizes reflect commonly used sample sizes in the field, future studies could have more statistical power through using larger samples. With the current study, we can not study all possible factors impacting on conformity to robots. For instance we do not know how the robots are perceived by the participants or how participants judge the visual acuity of the robots. Allen argued that a greater similarity between the participant and the confederates will increase the likelihood of the participant perceiving the confederates as an appropriate reference group and hence will increase the level of conformity (*44*). Thus, adults might not form social bonds with small humanoid robots, but only with larger adult-size robots. Children on the other hand might not want to disagree with the robots for reasons that are as yet unexplored. All properties of design and behavior of the robots might potentially be factors that produce an influence on social conformity which need to be explored in future research.

Materials and Methods

We followed the experimental procedure as outlined by Asch (*20–22*) and used the same stimulus specification where possible (*22*). The adult experiments took place within a university lab setting while the experiments with the children were conducted at a local primary school in an empty classroom. Rather than presenting the stimuli on card, a TV screen was used. In the robot-peer condition software remotely orchestrated the response behavior of the three robots via a wireless network. The confederates, both human and robot, all followed the same pat-

tern of responses. All responses from participants and confederates were reported vocally and recorded by the experimenter using pen and paper. Participants (and confederates) were seated around a table, facing the TV screen (Fig. 1, B and C). For each of the 18 trials (12 critical, 6 neutral) the experimenter recorded the responses in a clockwise direction, beginning with the confederates and finishing with the participant. This order was constant for the human-peer and robot-peer conditions as was the seating plan. In the control condition no confederates were present.

Participants

60 adults took part in the experiment: 28 males ($M_{age} = 30.32$ years, $SD = 13.76$) and 34 females ($M_{age} = 31.48$, $SD = 14.61$). Participants were recruited via the online subject pool maintained by the School of Psychology at the University of Plymouth and were paid £4. They were randomly assigned to one of three conditions (control, robot peer, human peer), none of the participants were excluded (exclusion criterion: not using required vision correction). As participants were recruited through volunteer sampling, based on our one-way balanced between subjects design with three groups, the sample had a power level of .78 to detect a medium to large effect ($f = 0.4$) assuming an alpha level of .05.

43 children took part: 21 boys ($M_{age} = 8.47$ years, $SD = 0.58$) and 22 girls ($M_{age} = 8.50$, $SD = 0.50$). All were pupils at a local primary school in the Plymouth (UK) area and consent was obtained from both the school and parents. Children were pooled from one of two classes: Year 3 (aged 7 to 8, $n = 21$) and Year 4 (aged 8 to 9, $n = 22$). We have selected this age group as it is well-studied with respect to conformity, cf. (*45*), and younger children might not understand the task, as suggested by (*29*). Children were randomly assigned to either the control or robot-peer condition. Children would be excluded if they were not using required vision correction or if they felt uncomfortable. No children were excluded. The experimental sessions took place

over the course of a single school day and were located within a spare classroom within the school. No reward was provided, however at the end of the day a small presentation about robots was given by the experimenter. A power analysis showed that we had $> .71$ power to detect a medium to large effect ($d = .8$) assuming an alpha level of .05.

Materials

The length and order of the target and comparison lines were identical to the specifications outlined in original Asch studies (*20, 22*), see Table S2. A 32 inch LCD TV was used to display the stimuli as opposed to physical cards with printed lines. A laptop was connected to the screen running custom software to display the stimuli. In the human-peer condition the laptop's screen, only visible to the first confederate, also displayed the confederate answer allowing the first confederate to read this while looking at the TV screen. In the robot-peer condition this software was also used to orchestrate the behavior of the robots over a WiFi network.

The use of the TV screen introduced a deviation from the original Asch setup. We were unable to separate the target line and the matching comparison line by 40 inches (101.6 cm) as the TV screen was not wide enough for this. Instead we held this distance between the target line and the left hand comparison line constant at 40 cm. The horizontal distance between the edge of the screen and target line/right hand comparison line was 8.3 cm. All other dimensions were in accordance with the original experiments (*22*), see also Fig. S1 and Table S3. A smaller separation of target line and comparison lines makes the stimuli less ambiguous as it permits an easier comparison of line lengths, which should have no implications in studying normative social influence.

Three SoftBank Robotics Nao humanoid robots (Fig. 1A) were used as the confederates in the robot-peer condition. The Nao is a small 25 degree-of-freedom 58cm tall humanoid robot designed primarily for human-robot interaction. Each robot was autonomous, running custom

software that allowed it to be controlled by the software running on the experimenter's laptop. This software performed scripted behaviors that were run each time a new trial was displayed. The robots were seated at the table. In Experiment 1 they were seated on plastic boxes to elevate their position relative to the adult subjects (see Fig. 1C) to obtain approximately the same difference in face height between participant and robots across experiments. Only power cables were connected to the robots. The robots' head motor joint positions required to gaze at the TV screen, experimenter and participant were preprogrammed.

Procedure

Experiment 1

Subjects were randomly assigned to one of the three following conditions. In the 'control' condition the participants completed the task on their own, providing a baseline measure of performance. In the 'human-peer' condition the participants completed the task with three human confederates, serving as a replication of the original Asch experiments. In the 'robot-peer' condition the human confederates were replaced by robots.

Upon arrival in the experiment room, the confederates sat down in their agreed positions ensuring that the participant sat in the last seat (Fig. 1C). Participants (including the confederates) were briefed and consent was received. In the robot-peer condition, the briefing and obtaining of consent took place prior to entering the room. The robot's were already seated around the table when the participant entered.

Each participant was presented with an information sheet and a consent form. Participants were informed on the information sheet that they needed to perform a simple visual discrimination task in which they needed to indicate which of three comparison lines matched the length of a standard line in 18 such comparisons. They were also informed that all answers would be recorded on a prepared form.

An example visual stimulus was then used to provide a tangible instruction of the task. Participants were then offered the opportunity to ask for clarifications. Except in the control conditions, the experimenter defined the order of responses, clock-wise beginning with the first confederate. Following this the experiment began.

In the control condition participants performed the task alone, with only the experimenter in the room. In the human-peer condition the confederates provided their responses first. The first confederate was located opposite the participant, allowing the first confederate to see the laptop screen displaying the confederate answer while gazing toward the TV screen. All the other confederates followed her response. All robot confederates provided their response first as well.

Debriefing took place immediately after the experiment finished. Participants in the control condition were informed that they were in a control condition for the experiment. The nature of the experiment was also explained to them. Participants in the human- and robot-peer conditions were informed of the role of the confederates and what the aim of the experiment was: the measuring of normative social conformity. They also were given a questionnaire to collect demographic details, data on familiarity with and views of robots, and a personality test. All participants were requested to maintain confidentiality to avoid biasing future experiments.

Experiment 2

Experiment 2 mainly followed the same experimental procedure as described for Experiment 1. In Experiment 2, child subjects were only subject to the control and robot-peer conditions to which they were randomly assigned. Children were briefed while sitting at the table in the experiment room. Parental consent was obtained in advance. The children were not given any information sheet or questionnaire. The experimenter informed them orally that they needed to perform an “eye test” in which they needed to indicate which of three comparison lines

matched the length of a standard line in 18 such comparisons. They were also informed that all answers would be recorded on a prepared form. From here on, the course of the experiment was exactly the same as for the robot peer and control condition of Experiment 1, including the practice trial, the opportunity to ask for clarifications, the order of responses, and debriefing in the control condition. In the robot-peer condition, children were told during debriefing that the robots were trying to "trick" them and see whether they would agree with the robots. Children were also asked not to tell others about the experiment to avoid biasing future experiments.

Presentation of the robots

In the conditions where robots acted as confederates, the robots did not react to the participant when they entered and sat down. The experimenter outlined the instructions for the visual discrimination task and provided an example of the visual stimuli. When the lines were shown on screen the robots all gazed toward the experimenter as if listening to the instructions. The presentation of the real experimental trials commenced after this. From this stage onward, the scripted behavior of the robots was initiated each time the experimenter used the laptop to display the next set of comparison lines on the TV screen: all robots were instructed to gaze towards the screen, each with a different motor speed randomly selected uniformly from a given range. The robots paused for a random period between 0.75 and 1.5 seconds and then verbalised the desired response via an on-board text-to-speech engine. After giving a response, a robot occasionally looked at the participant for 1.5 seconds and then looked back at the screen. The purpose of this gaze behavior is to apply a certain amount of social pressure on the participants. A flow diagram of the scripted robot behaviour during the experimental trials can be found in Fig. S2.

A large part of this experiment depended on the manner in which the confederates were presented to the participant, particularly in the case of the robots. As such, care was taken

to present and treat the robots as individual social entities through the observable behavior, and how they were treated by the experimenter (i.e. the behavior of the experimenter directed toward the robots).

To provide the robots with a basic level of animacy, each robot was programmed to exhibit small behaviors to avoid the robot appearing static. Small motor movements were executed around the given gaze direction as were movements of the wrist joints and fingers. These motor commands were executed at random within a given time frame. Blinking behavior was also introduced through toggling power to the LED eyes at random intervals. Each of the robots was provided with an individual voice through altering the pitch of the text-to-speech engine. The eye colour of each robot was also individual. Fiducial markers were placed in the four outer corners of the screen, to allow to robot to see the screen.

Throughout the experiments, the experimenter's behavior toward the robots was as similar as possible to their behavior toward the participant. For example, during the task description, eye contact was made with both the participant and each individual robot. The robots were also given and referred to by names: Snap, Crackle and Pop.

In the robot-peer condition, adult subjects were informed in the information sheet that the aim of the research is to investigate visual discrimination in humans and robots and that each experiment involved 4 participants (a mixed group of humans and robots). Other than this, the reasoning for the robots being present was kept unspecified.

Ethics

The research design for this study was reviewed and approved by the Plymouth University Ethics Committee for the Faculty of Science and Engineering. Adult participants provided informed consent prior to the experiment and informed consent was provided by the parents of children prior to the experiments. Full debriefing in all conditions took place immediately after

the experiment ended.

Supplementary Material

Analysis of Logit (Logistic Regression) model.

Fig. S1. Specifications of visual stimuli presented to the participants.

Fig. S2. Flow diagram of the scripted robot behavior during the experimental trials.

Table S1. Discrimination accuracy across conditions.

Table S2. Specification of standard and comparison line lengths.

Table S3. Dimensions of the stimuli presentation.

Data S1. Text file of adult participant responses in Experiment 1.

Data S2. Text file of child participant responses in Experiment 2.

Further data

References

1. C. Breazeal, Sociable machines: Expressive social exchange between humans and robots, Ph.D. thesis, MIT, Artificial Intelligence Laboratory (2000).

2. T. Kanda, T. Hirano, D. Eaton, H. Ishiguro, *Human-computer interaction* **19**, 61 (2004).

3. T. Belpaeme, *et al.*, *Journal of Human-Robot Interaction* **1**, 33 (2012). In press.

4. W. Burgard, *et al.*, *Artificial intelligence* **114**, 3 (1999).

5. K. Wada, T. Shibata, *IEEE Transactions on Robotics* **23**, 972 (2007).

6. P. H. Kahn Jr, B. Friedman, D. R. Pérez-Granados, N. G. Freier, *Interaction Studies* **7**, 405 (2006).

7. F. Tanaka, A. Cicourel, J. R. Movellan, *Proceedings of the National Academy of Sciences* **104**, 17954 (2007).

8. C. Breazeal, *et al.*, *Topics in cognitive science* (2016).

9. J. Kennedy, P. Baxter, E. Senft, T. Belpaeme, *2016 11th ACM/IEEE International Conference on Human-Robot Interaction (HRI)* (IEEE, 2016), pp. 231–238.

10. T. Fong, I. Nourbakhsh, K. Dautenhahn, *Robotics and autonomous systems* **42**, 143 (2003).

11. P. H. Kahn, H. E. Gary, S. Shen, *Child Development Perspectives* **7**, 32 (2013).

12. G. F. Melson, *et al.*, *Journal of Applied Developmental Psychology* **30**, 92 (2009).

13. E. S. Kim, *et al.*, *Journal of autism and developmental disorders* **43**, 1038 (2013).

14. B. Scassellati, H. Admoni, M. Mataric, *Annual review of biomedical engineering* **14**, 275 (2012).

15. S. Thill, C. A. Pop, T. Belpaeme, T. Ziemke, B. Vanderborght, *Paladyn, Journal of Behavioral Robotics* **3**, 209 (2012).

16. M. Deutsch, H. B. Gerard, *The journal of abnormal and social psychology* **51**, 629 (1955).

17. B. Reeves, C. Nass, *How people treat computers, television, and new media like real people and places* (CSLI Publications and Cambridge university press Cambridge, UK, 1996).

18. C. Nass, Y. Moon, *Journal of social issues* **56**, 81 (2000).

19. C. Nass, J. Steuer, E. R. Tauber, *Proceedings of the SIGCHI conference on Human factors in computing systems* (ACM, 1994), pp. 72–78.

20. S. E. Asch, *Groups, leadership, and men* pp. 222–236 (1951).

21. S. E. Asch, *Readings about the social animal* **193**, 17 (1955).

22. S. E. Asch, *Psychological monographs: General and applied* **70**, 1 (1956).

23. C. Nass, Y. Moon, P. Carney, *Journal of Applied Social Psychology* **29**, 1093 (1999).

24. C. Nass, Y. Moon, N. Green, *Journal of applied social psychology* **27**, 864 (1997).

25. D. Abrams, M. Wetherell, S. Cochrane, M. A. Hogg, J. C. Turner, *British Journal of Social Psychology* **29**, 97 (1990).

26. T. F. Linde, C. Patterson, *The Journal of Abnormal and Social Psychology* **68**, 115 (1964).

27. A. H. Eagly, *Psychological Bulletin* **85**, 86 (1978).

28. R. Bond, P. B. Smith, *Psychological bulletin* **119**, 111 (1996).

29. M. B. Walker, M. G. Andrade, *The Journal of social psychology* **136**, 367 (1996).

30. M. Pasupathi, *Psychology and Aging* **14**, 170 (1999).

31. P. R. Costanzo, M. E. Shaw, *Child development* pp. 967–975 (1966).

32. T. J. Berndt, *Developmental psychology* **15**, 608 (1979).

33. J. Brandstetter, *et al.*, *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems* (IEEE, 2014), pp. 1335–1340.

34. D. Haun, M. Tomasello, *Child development* **82**, 1759 (2011).

35. A. Hanayama, K. Mori, *Psychology* **2**, 661 (2011).

36. R. M. Bond, *et al.*, *Nature* **489**, 295 (2012).

37. A. D. Kramer, J. E. Guillory, J. T. Hancock, *Proceedings of the National Academy of Sciences* **111**, 8788 (2014).

38. N. Salomons, M. van der Linden, S. Strohkorb Sebo, B. Scassellati, *Proceedings of the 2018 ACM/IEEE International Conference on Human-Robot Interaction* (ACM, 2018), pp. 187–195.

39. N. Sharkey, A. Sharkey, *Interaction Studies* **11**, 161 (2010).

40. A. Caliskan, J. J. Bryson, A. Narayanan, *Science* **356**, 183 (2017).

41. J. Fasola, M. J. Mataric, *Proceedings of the IEEE* **100**, 2512 (2012).

42. J.-H. Han, M.-H. Jo, V. Jones, J.-H. Jo, *Journal of Information Processing Systems* **4**, 159 (2008).

43. K. Mori, M. Arai, *International Journal of Psychology* **45**, 390 (2010).

44. V. L. Allen, *Advances in experimental social psychology* (Elsevier, 1965), vol. 2, pp. 133–175.

45. D. B. Haun, van Leeuwen Edwin J.C., M. G. Edelson, *Developmental Cognitive Neuroscience* **3**, 61 (2013).

Acknowledgements: We thank Matthew Rule for acting as the human third confederate in the human-peer condition, Emily Ashurst for assisting with running the child experiments, and Paul Baxter for his insightful comments. **Funding:** This work was funded by the EU FP7 ALIZ-E (248116), FP7 DREAM (611391), FP7 Marie Curie Actions ITN RobotDoC (235065) and H2020 L2TOR (688014) projects, and grants from the Cluster of Excellence Cognitive Interaction Technology 'CITEC' (EXC 277), Bielefeld University. **Author contributions:** DT

and ALV conceived the initial experiment with adults. DT, ALV, RR and TB designed and planned the experiments. ALV secured ethical approval. RR designed and developed the software for the laptops and robots. DT, ALV, and RR conducted the adult experiments where DT was the experimenter in all conditions. ALV and RR were confederates in the adult human-peer conditions. RR conducted the experiments with the school children. DT, RR and TB performed the data analysis. All authors contributed to the paper. **Competing interests:** The authors declare no competing financial interests. **Data and materials availability** All data needed to evaluate the conclusions in the paper are present in the paper or the Supplementary Materials.

Fig. 1. Overview of the experimental setup and visual stimulus. (**A**) The SoftBank Robotics Nao humanoid robot used as confederate. (**B**) Overview of the participant seating arrangement. In the control condition only the participant and experimenter were present. Participants' judgments are collected in a clockwise order beginning with the confederates and ending with the subject. (**C**) Illustration of the arrangement in a real setup. (**D**) Illustration of the visual stimuli presented to participants via a computer screen. The target line is located on the left and the three labeled comparison lines are located on the right. Participants say which of these matches the length of the target line.

Fig. 2. Discrimination accuracy across conditions. (**A**) The mean accuracy of the adults for the critical and neutral trials, across each experimental condition (control $n = 20$, robot peer $n = 20$, human peer $n = 20$). During the critical trials the presence of human peers leads to a significant decrease in discrimination accuracy due to subjects conforming with the human confederates. (**B**) The mean accuracy of the children during the discrimination task (control $n = 21$, robot peer $n = 22$, no human-peer condition). During the critical trials the presence of the robot-peers lead to a significant decrease in accuracy due to group conformity. Error bars denote 95% Confidence interval of the mean estimate; likelihood ratio test on logistic regression, * $P < .01$; ** $P < .001$.

Fig. 3. Breakdown of incorrect participant responses. The bars shows the ratio of conforming (i.e. going with the confederates' response) against non-conforming responses in the critical trials; for the adults in the human-peer condition ($n = 20$) and for the children in the robot-peer condition ($n = 22$). 83% of all incorrect responses from the adults were found to be conforming with the group of human confederates while children's conformity with the robots was 74%. Two-tailed χ^2 test, ** $P < .001$.