Humans Conform to Robots: Disambiguating Trust, Truth, and Conformity

Nicole Salomons, Michael van der Linden, Sarah Strokhorb Sebo, Brian Scassellati
Department of Computer Science
Yale University
{nicole.salomons,michael.van.der.linden,sarah.sebo,brian.scassellati}@yale.edu

ABSTRACT

Asch's [2] conformity experiment has shown that people are prone to adjusting their view to match those of group members even when they believe the answer of the group to be wrong. Previous studies have attempted to replicate Asch's experiment with a group of robots but have failed to observe conformity [7, 25]. One explanation can be made using Hodges and Geyers work [17], in which they propose that people consider distinct criteria (truth, trust, and social solidarity) when deciding whether to conform to others. In order to study how trust and truth affect conformity, we propose an experiment in which participants play a game with three robots, in which there are no objective answers. We measured how many times participants changed their preliminary answers to match the group of robots' in their final answer. We conducted a between-subjects study ($N = 30$) in which there were two conditions: one in which participants saw the group of robots' preliminary answer before deciding their final answer, and a control condition in which they did not know the robots' preliminary answer. Participants in the experimental condition conformed significantly more (29%) than participants in the control condition (6%). Therefore we have shown that groups of robots can cause people to conform to them. Additionally trust plays a role in conformity: initially, participants conformed to robots at a similar rate to Asch's participants, however, many participants stop conforming later in the game when trust is lost due to the robots choosing an incorrect answer.

CCS CONCEPTS

• **Human-centered computing** → **User studies**; • **Computer systems organization** → *Robotics*;

KEYWORDS

Conformity; Human-Robot Groups; Trust

ACM Reference format:
Nicole Salomons, Michael van der Linden, Sarah Strokhorb Sebo, Brian Scassellati. 2018. Humans Conform to Robots: Disambiguating Trust, Truth, and Conformity. In *Proceedings of HRI'18, Chicago, IL, USA, March 5th–8, 2018,* 9 pages.
https://doi.org/10.1145/3171221.3171282

HRI'18, March 5th–8, 2018, Chicago, IL, USA

ACM ISBN 978-1-4503-4953-6/18/03…$15.00
https://doi.org/10.1145/3171221.3171282

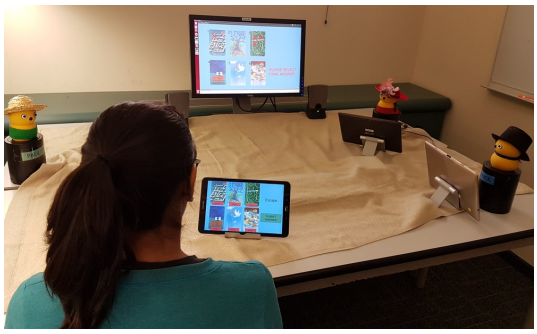

Figure 1: In this experiment, participants play a game with three MyKeepon robots. In the game, participants are given opportunities to change their answer, and often show conformity to the group of robots.

1 INTRODUCTION

"Conformity refers to the act of changing one's behavior to match the responses of others" (page 606)[9]. One of the most foundational psychological studies measuring conformity was performed by Asch in the 1950's [2]. When individually asked to answer a very simple perceptual question (identifying which line out of a set of lines matched another line in length), participants answered correctly 99% of the time. However, when placed second to last in a group of confederates, if the confederates unanimously verbalized an incorrect answer before them, the participant would choose the same incorrect answer 37% of the time.

Shiomi et al. and Brandstetter et al. attempted to replicate the Asch paradigm with robot confederates instead of humans [7, 25], but neither study was able to show conformity to a group of robots. Possible reasons suggested for the lack of conformity included lack of a social relationship between the participant and robots, participants not viewing the robots as authoritative entities, and the robots not being human-like enough [7, 25].

These results are somewhat surprising, as research has shown that individual robots are capable of persuading people. Siegel et al. showed that participants voluntarily donated money to a persuasive robot [27]. Another example of persuasive robots can be seen in research conducted by Chidambaram et al., in which participants comply with a robot's suggestions when playing a game [8]. Additionally, in cases such as lexical entrainment, the influence of robots can persist even after the interaction ends [6]. It is therefore unexpected that groups of robots fail to provoke conformity in their interactions with people.

Hodges and Geyer [17] offer one possible alternative interpretation of Asch's conformity experiment, in which they propose that participants were constrained by multiple influencing factors including trust, truth, and conformity. They argue that participants were placed in a challenging situation in which they had to trade-off their trust in the group members, their desire to give a truthful answer, and the pressure to conform to the group. They support this argument by pointing out the time-varying nature of participant responses, often interweaving correct (non-conforming) answers with false (conforming) answers, and the overall low conformity rate (37%).

Robots cause less social pressure then humans and their level of trustworthiness is unclear when compared to a person. We propose to unravel the interplay between trust, truth and conformity in human-robot groups by changing the task to one where there is no obvious objective truth, and therefore analyze the effect of trust in the interaction.

2 BACKGROUND

In this section we will review relevant literature on conformity research in psychology, conformity with robots, and research conducted with multiple social robots.

2.1 Conformity

Conformity is defined as changes in one's behavior to match the behavior of the rest of a group [9]. Asch conducted a series of experiments [2], in which he showed that humans conform to the answers of a group. In the experiments, participants would complete a visual task where, given a picture of a line, they were asked to choose a line of equal length out of three options (Figure 2). Out of 18 rounds, 12 were critical rounds, in which six to eight actors would say the wrong answer out loud and afterwards the participant would choose an answer. The correct answer was fairly obvious, and when working alone, participants answered correctly 99% of time. However, when participants were placed in a group of actors who all chose the equivalent wrong answer, the participant would conform to the group, answering incorrectly in 37% of the rounds on average. Asch conducted additional experiments [2] varying the number of trials, the number of actors, and the ambiguity of the lines. He concluded that in all the different variations, people consistently conformed to the group. Asch believed that people disregard answering questions correctly in order to answer in accordance with a group due to the perceived social pressure participants feel from the actors. This made them not want to answer differently from the rest of the group.

Asch's participants were likely acting due to peer pressure. But other reasons for conforming to a group exist. Deutsch and Gerard suggest two main reasons as to why people conform to a group: normative conformity and informational conformity [11]. Normative conformity pertains to peer pressure, or the conformation to the expectations of others. Informational conformity is conformity due to uncertainty about one's own answer or behavior.

2.2 Human Conformity in Robotics

Similar experiments were conducted by Beckner et al. and Brandstetter et al. where conformity was tested with both a group of robots and a group of human actors [4, 7]. In addition to the traditional Asch conformity line test, they also tested verbal tasks (determining verb tenses). In both tasks, four NAO Aldebaran robots all stated the same wrong answer in certain rounds. Although the study did show conformity with human actors, it did not show conformity with a group of robots. It also did not achieve significant differences between the robot condition and a baseline (participants would press a button instead of saying it out loud).

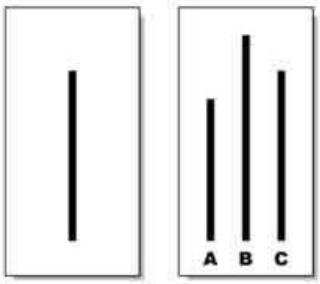

Figure 2: In Asch's conformity experiment, participants were asked to match the length of the left line to the right lines [22].

Shiomi and Hagita also tested conformity with robots [25]. Using Asch's line task with groups of robots, in which 12 out of the 18 rounds they measured participant's decision to conform to the robots. They had two different conditions: one in which the robots all synchronized their actions, looking first at the previous robot and then at the next robot before giving the answer, and a second condition where the robots did not synchronize their actions. Neither approach resulted in significant conformity when comparing to an alone condition where the participant was the only person to respond to the line comparison task.

People's tendency to conform to humans did not extend to robots in these previous studies. However research has shown that robots are capable of changing people's behavior using persuasion. Robots have been able to persuade people using verbal comments, body language and gaze [8, 14]. Examples of robots persuading people can be seen in education [5], in motivating people [23], in health self-management [21], in energy conservation [15] and in advertisement [26]. Similar to persuasion, conformity causes someone to change their behavior but without showing direct intent to do so. Singular robots can cause persuasion, but it is still unclear if multiple robots cause conformity.

2.3 Groups of Robots

Preliminary work has been conducted focusing on how people engage and perceive groups of robots. Fraune et al. analyzed how people respond to videos of single versus groups of NAO robots, iCreate robots and Pleo robots (which are dinosaur-like robots) [13]. Their results show that participants had more negative responses (such as feeling more threat, anxiety and lack of trust) to groups of iCreate robots compared to individual iCreate robots. But on the other hand, when interacting with a more anthropomorphic robot, participants had more positive responses to groups of NAOs rather than individual NAOs. There wasn't a large difference in response when interacting with individual or groups of Pleos.

Admoni et al. compared perceptions of robot members belonging to either minority or majority groups of robots[1]. Participants evaluated videos of robots which would either be dancing along with the rest of the group, or would be dancing differently to the rest of the group. Robots that were dancing differently than the group were rated as more creative, more anti-social and less cooperative, compared to robots who danced along with the group. Demonstrating a similar result, Fraune et al. showed that people preferred groups of diverse robots over homogeneous groups of robots, and found the diverse groups less threatening [12].

Other studies involving multiple social robots were conducted on teaching emotional understanding to children [19] and [20]. Nawroj et al. studied the effects of mimicry, appearance and gaze in establishing social groups of robots [24]. Limited research so far has been conducted on evaluating how people interact with groups of social robots. In this paper, we aim to explore the effect that multiple robots cause on human behavior, and whether they are capable of changing people's behavior.

3 METHODOLOGY

Hodges and Geyer suggest that situations such as the one presented to participants in Asch's experiments made participants need to trade off on truth, trust and social solidarity [16, 17]. The truth aspect pertains to wanting to maintain a view that is accurate to what one believes, trust pertains to how reliable the other sources are believed to be and social solidarity pertains to incorporation of both one's own view and other people's views when choosing answers. In this study we focus only on trust and truth and how they relate to conformity, as we believe social solidarity to be more relevant when the group needs to reach consensus such as collaboration scenarios.

It is likely that the previous studies with robots did not show conformity because the tasks had a clear correct answer. Thus when the robots chose an obviously false answer, the truth aspect overpowered the values of trust and social solidarity. This study separates the truth aspect so we have a clearer understanding of how the role of trust effects conformity.

In our experiment participants play a game with three robots in which there is no obvious correct answer. This means that prior knowledge is not a factor, and the answer is inherently subjective. We are also using a two-stage round in which the participant gives an initial and then final answer. This mechanic permits us to directly measure when they change their mind, and allows us to manipulate the robot's answers depending on the participants answer in order to maintain consistency across participants. A game with answers that are unknown to the participant beforehand is more likely to induce conformity than past studies. Additionally, we are interested in how the performance history of the robots effect the decision of the participants when deciding to conform or not. Does the decision of conforming to the robots stay consistent during the interaction, or does the performance of the robots determine whether participants conform?

There are two hypotheses for this study:

Hypothesis 1: *When provided with the opportunity to conform without providing an obviously false answer, participants will conform with a group of robots at a rate similar to the original Asch participants.*

Hypothesis 2: *When the performance of a group of robots degrades, participants conform less often to the the answers of the robots.*

In this study, we do a between-subjects human subject experiment in which there are two conditions:

- **Experimental Condition** - Participants in this condition can see the robots' initial answer in anonymous form. That is, they can can see how many robots chose a certain answer, but not exactly who chose which answer. After participants select their final answer, they see everyones individual final answer.
- **Control Condition** - Participants do not see the initial answers of the robots. They will only see everyones final answers.

The second condition is a baseline condition, in which we are testing how often participants change their answers without extra information, and also whether the robot's actions or utterances are causing them to change their answers.

3.1 Participants

30 participants were recruited of which 11 were male, and 19 were female, with average age 27 years old (SD = 9.6). Most of the participants were students from a local university and people from its surrounding town. Participants were randomly distributed between conditions: 15 participants (5 male, 10 female) were in the experimental condition, and 15 participants (6 male, 9 female) were in the control condition. They signed a consent form agreeing to participate in the study, and received five dollars compensation for their time. The game with the robots and the questionnaire took approximately 30 minutes.

3.2 Procedure

3.2.1 Cards. Participants played a modified version of the card game "Dixit" with the robots. Instead of using physical Dixit cards, they were digitalized and were shown on the participants tablet (examples of Dixit cards used in this experiment can be seen in Figure 3). In each round of the game, the participant and robots are shown six cards, each depicting a whimsical drawing.

3.2.2 The sequence of the task. At the beginning of the session the participant played a tutorial round with the robots, in which the game was explained in detail. The experimenter then left the room and the participant played 20 rounds of the game with the group of robots. The sequence of what happens in each round can be seen in Figure 3. Each round proceeded as follows:

(1) A word is given to the participant.
(2) The participant chooses the card that best represents the word out of six cards.
(3) In the experimental condition, how many times each card was selected is shown on a shared screen. In the control

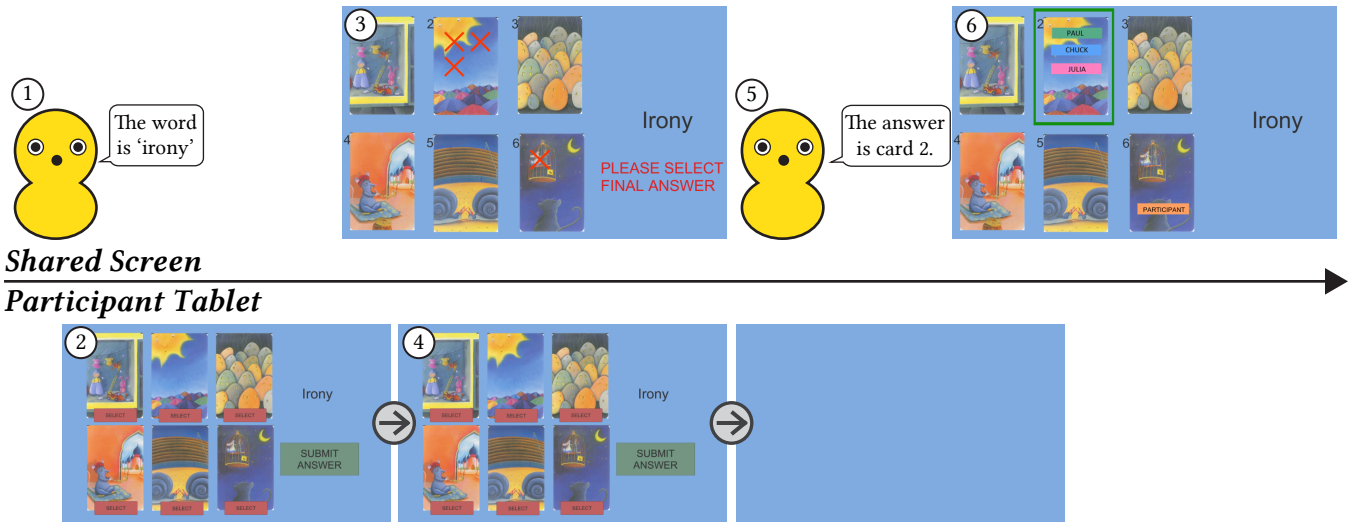

Figure 3: (1) A game master robot gives a word, (2) the participant chooses the card that best corresponds the word on their tablets, (3) the robots and participants anonymous answers are shown in the form of red "X's", (4) the participant chooses their final answer, (5) the game master gives the correct answer, (6) the robots and participants final answers are shown by having their name on top of the card they chose on the shared screen.

condition, only the cards are shown on the screen. The differences between the conditions can be see in Figure 4.

(4) The participant may choose to change their answer to one of the other six cards (for the same word).
(5) The correct answer is given.
(6) The robot's and participant's answers are shown publicly to everyone.

After the participant has played 20 rounds with the robots, they complete a post-experiment questionnaire which is detailed in Section 3.5.

3.2.3 User Interfaces: Tablet and Shared Screen. There were two different user interfaces used in the study: the tablet user interface, and the shared screen. The participant and each of the robots have their own tablet on which they choose their preliminary and final answers for the current word (Figure 3 - (2) and (4)). When the participant is not selecting an answer, the tablet will have a blank blue screen.

There is also a shared interface which is visible to the participant and to all the robots, which we call the shared screen. The shared screen shows videos of a fourth robot called the **game master**, which gives the word at the start of each round and will also provide the correct answers at the end of each round (Figure 3 - (1) and (5)). Additionally the shared screen shows the anonymous answers in the experimental round (Figure 3 - (3)), and each of their final answers in both conditions (Figure 3 - (6)).

The shared screen facilitates social pressure on the participant as the participant and the robots will all be looking at the same screen. The participant will feel as if the robots clearly can see when the participant chose a different answer than them, thus participants may feel peer pressure from the group.

3.2.4 Preliminary Answers and Final Answers. The difference in the two conditions resides in how the preliminary answer of the participant and robot is shown. In the experimental condition, each of the six cards is shown on the shared screen and red "X's" indicate how many people/robots chose that card. For example in Figure 4a, three people/robots chose one card, and the other person/robot chose a different card, but no information is given about which red "X" corresponds to whom. In the control condition, no information is given about each of the preliminary answers, instead they are just shown a screen with the cards (Figure 4b) and are asked to select a final answer.

After the participant selects their final answer, all of the answers will be publicly shown on the shared screen. The names of the robots and the name of the participant will appear on top of the card they chose. Both the control and the experimental condition will see everyone's final answers.

3.3 Rounds

Participants played a total of 20 rounds with the robots, in which the word and the specific cards for the round were chosen beforehand. There were varying types of rounds to make the game feel realistic. There are rounds when robots change their answers to show that it is permitted to change their answer; likewise, there are rounds where robots do not change their answers, to show that it is also permissible to choose their original answer. Table 1 shows the type of each round. The rounds corresponded to the following types:

- **Unanimous Rounds** - The answer is fairly obvious, and all the robots choose the same answer in both the preliminary and final round. The participant is also expected to choose this answer.
- **One Robot Converges** - Two robots choose one answer (usually the same as the participant) and the third robot

(a) Preliminary answers - Experimental Condition

(b) Preliminary answers - Control Condition

Figure 4: The preliminary answers of the robots in the experimental and control conditions. In the experimental it is visible to the participant exactly which cards are chosen and how many robots chose each card in the form of red "X's" on top of the cards. In the control condition, no information is given about which cards were chosen.

chooses a different answer in the preliminary round and converges to match the rest of the group for its final answer.

- **Two Robots Converge** - One robot chooses the same answer as the participant and the other two choose a different answer. The two differing robots change to conform with the participant and other robot.
- **One Robot different** - One robot is different from the three others (two robots and participant) and it does not change its answer.
- **All different** - All the robots choose a different answer from each other. In some of the rounds one of the robots might overlap with the participant's answer.
- **Critical Rounds** - All the robots choose the same opposing answer from the participant. These are the rounds in which we are testing conformity.

Out of the 20 rounds, six are **critical rounds**. In these rounds, the three robots are programmed to unanimously choose a different reasonable answer opposing the participant's answer. For example: if both cards one and two are reasonable answers for a word - when the participant chooses card one, then all the robots will choose card two, and vice versa. The six critical rounds where the robots unanimously diverge their answers, are the rounds in which we will observe whether people conform to robots and can be seen in Table 2. In three of the six rounds the participants initial guess would be right, and in three of the rounds, the robot's preliminary guess would be correct. We keep this balanced to prevent participants from believing that the robots are always correct or always incorrect.

Round Numbers	Type of Rounds
1,3,6,8,14,18	Unanimous Round
2,4	One Robot converges
11,19	Two Robots Converge
12	One robot different
7,15,17	All different
5, 9, 10, 13, 16, 20	Critical Round

Table 1: The round number with their type of round.

3.4 Keepon Robots

As shown in Figure 5, we used three MyKeepon robots, the consumer-grade version of a research robot called Keepon Pro [18]. The MyKeepon robot is a 15cm snowman-shaped robot with a foam exterior which was manufactured originally as a toy for young children but then modified by researchers for greater control and configuration. It has a flexible body with four degrees of freedom (pan, roll, tilt, bop). Each of the robots has a unique name, a different recorded voice and is dressed differently so they appear to have different personalities. During the rounds in which we are testing for conformity, the robots do not say anything to avoid a confound by direct verbal persuasion by the robots.

We chose three robots as a group size, because three people/robots are the minimum number of individuals usually considered to form a group. Asch found that increasing group size increased conformity, but found few differences when the group exceeded three confederates [2]. Of the four robots used in the study (three robots playing the game and one game master robot), we arbitrarily set the game master to be female and balanced the genders of the three remaining robots (two male and one female). The robots are seated around the table with the participant, and each robot has their own tablet to appear as they are playing the game with the participant.

Figure 5: In this human subjects study, a human participant interacts with a group of MyKeepon robots. Each of the robots is dressed uniquely and has a different voice.

Round	Word	Correct	Possible Solution Images
5	Escape	Robots	
9	Irony	Robots	
10	Beauty	Human	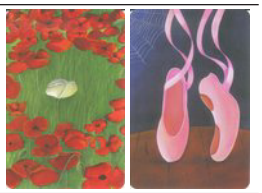
13	Immense	Human	
16	Control	Robots	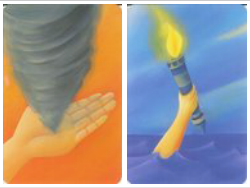
20	Lonely	Human	

Table 2: The critical rounds numbers, with their words, who was correct in the preliminary rounds and their two reasonable picture answers.

3.5 Measures

Our two main types of measures were: critical round data, which is the data collected automatically during the interaction with the robots; and questionnaire data which participants answered after playing the game with the robots.

3.5.1 Critical Rounds. Our main measurement is how often participants change their answers to the answer of the group of robots in the six critical rounds (in these rounds we measure conformity). Additionally, we measured how often participants continue conforming to the group in the next critical round, depending on if they got the answer right or wrong when conforming in the previous critical round.

3.5.2 Questionnaire. After playing the game with robots, the participants completed a post-experiment survey. The survey included part of the Godspeed questionnaire [3] with questions that assessed the robots' animacy, intelligence, and likability on a 5-point Likert scale. The survey also asked the participants to rate the following on a Likert scale from 1 (agree) to 5 (disagree): 1) the robots are better at playing this game than me and 2) I felt pressure to change my answers because of the robots.

The last question asked for an open-ended response: "Did you ever change your answer because of the robot and why?" We coded the open ended question answers into four different categories:

- **No** - The participant stated that they have never changed their answer because of the robots.
- **Peer pressure** - The participants says they changed their answer due to perceived peer pressure, or not wishing to be different from the rest of the group.
- **Robot Intelligence** - They changed their answer either because they believe the robots to be smarter than them, or because they mention the group being better than them individually.
- **Other** - They stated other reasons.

The data was categorized into these different types of answers, for us to be able to distinguish if participants were more prone to informational conformity or to normative conformity.

4 RESULTS

Our main goal was to analyze whether the experimental group conformed significantly more than the control group, that is: people who can see the robots' anonymous answers will change their answers to match the robots significantly more than when they can not see the robots anonymous answers.

4.1 Human-Robot Conformity

4.1.1 Critical Rounds. The main metric was whether participants conformed to the consensus of the group of robots during the six critical rounds. On average participants conformed a higher percentage of the time in the experimental group ($M = 28.9\%, SD = 23.9\%$) compared to the control group ($M = 5.6\%, SD = 12.1\%$). The difference was statistically significant using a mixed model Analysis of Variance (ANOVA) test ($p = 0.002$). The results can be seen in Figure 6. Additionally 66% of participants in the experimental condition changed their answers to match the robots at least once. Only 20% of participants in the control condition changed their answers in critical rounds at least once. In the experimental condition 46% of changes were made in critical rounds versus 16% in the control condition.

In these results we see strong indications that people do conform their answers to match those of the robots in the critical rounds. They change their answers significantly more when knowing the robots' answers, than when they do not know the answers. Asch's participants conformed on average in 37% of the rounds, and participants in our study conform on average in 29% of the rounds. Therefore we have evidence to support that **Hypothesis 1** is true:

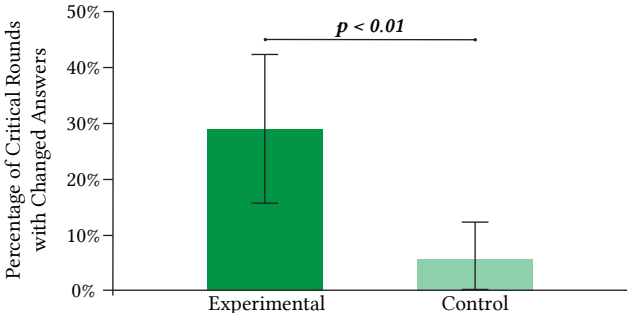

Figure 6: This graph shows the percentage of critical rounds with changed answers. Participants in the experimental condition changed their answers 29% of critical rounds, and participants in the control condition changed in 6% of the critical rounds.

when provided with the opportunity to conform without providing an obviously false answer, participants will conform with a group of robots. Our rate of conformity (29%) is similar, though not identical, to the rate Asch observed (37%).

4.1.2 Questionnaire. There was no significant difference between conditions in the ratings of animacy, likeability and intelligence that participants attributed to the robots. There were also no significant differences in conformity when comparing genders of the participants (p = .327). However, participants in the experimental condition reported feeling more pressure by the robots ($M = 2.73, SD = 1.334$), compared to the control condition ($M = 1.6, SD = 0.828$). The difference was significant ($p = .009$). There was also a correlation between reporting pressure to change and critical round changes. Using Pearson Correlation, they were moderately positively correlated ($R = 0.6305, N = 15, p = 0.012$).

The last question asked if they ever changed their answers because of the robots and why. 15% of participants in the experimental condition report that they were not influenced by the robots, and the remaining 85% responded positively to changing their answers because of the robots. In the control condition, 33% changed because of the robots, and the remaining 67% did not report changing because of the robots. For example one participant wrote that "Once (I changed my answer) because I wasn't sure and they all seemed sure". Another participant wrote: "Yes, because I was on the fence and their answers made me reconsider. I assumed I wouldn't be the lone right one". Another one wrote: "Yes, from life experiences, majority is usually correct". On the other hand, one participant did report purposefully going against the robots: "I once or twice changed my answers to go against the robots because I was feeling manipulated".

4.2 Trust and Conformity

When a participant in the experimental condition conformed to the robot group, and later found out that the robots were wrong, they rarely conformed to the robots again. Table 3 shows whether participants conformed to the robots in the following round depending on if they got the answer correct and if they conformed in the prior round. As the table shows, when the participant does not change their answer in the prior round, they will usually continue

	conformed	did not conform
changed answer - right	55%	45%
changed answer - wrong	29%	71%
not change - right	25%	75%
not change - wrong	26%	74%

Table 3: Percentage of time participants in the experimental condition conformed in the subsequent round depending on if they changed in the prior round and if they got the question right or wrong.

	No	Peer Pressure	Intelligence	Other
Experimental	15%	0%	69%	15%
Control	67%	0%	13%	20%

Table 4: This table displays the reasons why participants changed their answers in the different conditions.

not changing their answers in the following round independent of getting it right or wrong. When conforming with the robots, and getting it right, they are a little more likely to conform also in the next round. But once they conform, and get it wrong, they will be unlikely to conform in the next round. Additionally 56% of participants which conform and get it wrong, never conform to the robots again.

This is also supported by the questionnaire responses. Multiple participants mentioned that they initially decided to conform to the robots, but then stopped conforming when they saw that the robot did not achieve the correct answer. For example, one participant wrote: "Initially, because I thought robots would be good (at this task) and I went with the majority view. Later I completely disregarded the robots answers as they were mostly wrong anyway."

This result shows that we have evidence that indicates **Hypothesis 2** is true. Upon observing the robots fail by selecting the incorrect answer, the participants lower their conformity rate. Additionally the majority of participants never conform to a robot's answer again after the robot made a failure.

5 DISCUSSION

In this section we will discuss our measures of conformity, the reasons participants changed their answers, and show the role that truth, trust and conformity play.

5.1 Is this really conformity?

One of the main questions that arises is if this can really be considered conformity? The reason this question emerges is because different psychologists have different definitions for conformity. If we use Cialdini and Goldstein's [9] definition, then people do indeed conform to the robots because participants changed their responses to the group response on several occasions. Other psychologists define conformity as yielding to the pressure of a group [10], but in our study we do not have definite proof that participants are acting due to group pressure. In this paper we have used the former definition of conformity which states that conformity is when one changes their behavior to match the groups behavior.

5.2 Why do people change their answers?

Deutsch and Gerard propose two different reasons why people conform: normative conformity which is due to peer pressure, and informational conformity in which people conform because they are uncertain themselves [11]. In Asch's experiment the most likely reason participants were conforming to the group was due to peer pressure because in the rounds participants conformed, they were aware the answer was incorrect and answered incorrectly nonetheless to be part of the group. In this section we give an analysis studying whether normative conformity or informational conformity plays a larger role in people conforming to robots in the scenario of playing a game with no absolute correct answer.

One indicator of informational conformity was observed on critical round 16, in which none of the participants changed their answers to match the robots. The word was control and all of the participants chose the same answer (the card with a hand holding a tornado). This shows that when the question was less ambiguous, the participants would refuse to conform to the robots. This is an indicator peer pressure alone would not cause robot conformity. However the participants did report pressure to change in their questionnaire responses. As seen in the results, there was a positive correlation between critical round changes and reporting feeling pressured by the robots. Therefore this indicated that participants did feel some amount of peer pressure.

On the questionnaire we asked each participant an open ended question: "Did you ever change your answers because of the robots and why?". We have coded each participants answer into different categories: No (didn't report changing answer due to robots), due to peer pressure, due to robot intelligence, and others. As can be seen in Table 4, the majority of participants in the control group stated that they never changed their answers because of the robots (67%). Participants in the experimental condition however often stated that changed their answer either because they believe in strength of numbers or because they believed the robots were smarter than them (69%). 15% stated the robots did not influence their answers. But no participants reported feeling pressured by the robots.

In our study both peer pressure and robot group intelligence appear to play a role in participants changing their answers to the groups answer. But evidence (not conforming on a less ambiguous round and few participants reporting feeling pressured by the robots) suggests that most of the participants were acting on informational conformity, that is: conforming because they group provided information when they were unsure of their own answers.

5.3 Interplay between truth, trust and conformity

The relationships between truth, trust, and conformity are central to our analysis. Truth pertains to an individual's degree of confidence in their answer and their desire to give a correct answer. Trust, on the other hand, we define to be how confident an individual is in the abilities of the other members of the group. Trust is especially relevant to informational conformity (conformity due to uncertainty in owns answer). The level of expertise other members have in the topic heavily influences in your decision to conform, because you trust they have more knowledge in the answer than yourself.

In Asch's experiment, participants had both very high confidence in their own answers (truth) since it was a simple line length comparison task and a high confidence in the abilities of their cohort (trust) since they were all students of the same university. Therefore when the cohort chose an obvious false answer, the values of trust and truth came into conflict as they were aware that the answer of the group was incorrect, and it resulted in participants switching between the true answer and conforming with the group. Previous robotic studies which followed Asch's protocol placed participants in a situation where they had the same confidence in their own answers (high truth) but less certainty in the abilities of their robot partners (low trust). We propose the reason that participants did not conform to the group of robots was due to the combination of an uncertain level of trust in the robots and the obvious incorrect answer chosen by the robots.

In our study, we chose a task which has no known truth value, which allowed us a better understanding of how the role of trust effects the interaction. At the beginning of the session, participants were likely to be unsure of how much trust to attribute to the robots. We propose that the union of the unclear levels of trust attributed to the robots, in addition to a task without a clear answer, were the cause of a moderate amount of conformity of the participants in our study. Additionally we believe that when the trust in the robots was lost (upon seeing the robots choose an incorrect answer) it caused the participants to mostly stop conforming to the robots. In summary, this study proposes that when given a task where the correct answer is not known beforehand, the conformity rate depends on the level of trust attributed to the robots.

6 CONCLUSION

Previous studies have shown that people will change their answers or behaviors to match those of a group. Asch's conformity experiment [2] showed that people conform even when they disagree with the correct answer. Our primary objective was determining if this extends to groups of robots. Previous research did not show conformity with groups of robots when playing a game with an obvious correct answer [7, 25]. Therefore we tested if participants conform to robots when playing a game where there is no objective correct answer.

In this paper we have shown that people will change their answers to match those of the robots in around a third of the rounds of a our game, and the result is statistically significant compared to a baseline. The reason they conform is likely due to informational conformity: they believe the robots may be better at the task than them. But once participants lose trust in the robots, they also cease to conform to the group of robots.

7 ACKNOWLEDGEMENTS

This work was supported by the National Science Foundation award number 1139078. We would like to thank all the participants of our study for their time. Additionally we also want to thank David Villareal for help in creating the robots behaviors, and Meiying Qin and Alireza Taheri for help in designing the protocol of the game.

REFERENCES

[1] Henny Admoni, Bradley Hayes, David Feil-Seifer, Daniel Ullman, and Brian Scassellati. 2013. Dancing With Myself: The effect of majority group size on perceptions of majority and minority robot group members. In *CogSci*. 1708 – 1713.

[2] Solomon E Asch. 1956. Studies of independence and conformity: I. A minority of one against a unanimous majority. *Psychological monographs: General and applied* 70, 9 (1956), 1.

[3] Christoph Bartneck, Elizabeth Croft, and Dana Kulic. 2009. Measurement instruments for the anthropomorphism, animacy, likeability, perceived intelligence, and perceived safety of robots. *International Journal of Social Robotics* 1, 1 (2009), 71–81. DOI: http://dx.doi.org/10.1007/s12369-008-0001-3

[4] Clay Beckner, Péter Rácz, Jennifer Hay, Jürgen Brandstetter, and Christoph Bartneck. 2016. Participants conform to humans but not to humanoid robots in an English past tense formation task. *Journal of Language and Social Psychology* 35, 2 (2016), 158–179.

[5] Lykke Brogaard Bertel and Glenda Hannibal. 2016. Tema 2: The NAO robot as a Persuasive Educational and Entertainment Robot (PEER)–a case study on children's articulation, categorization and interaction with a social robot for learning. *Tidsskriftet Læring og Medier (LOM)* 8, 14 (2016).

[6] Jürgen Brandstetter, Clay Beckner, Eduardo Benitez Sandoval, and Christoph Bartneck. 2017. Persistent Lexical Entrainment in HRI. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 63–72.

[7] Jürgen Brandstetter, Péter Rácz, Clay Beckner, Eduardo B Sandoval, Jennifer Hay, and Christoph Bartneck. 2014. A peer pressure experiment: Recreation of the Asch conformity experiment with robots. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. 1335–1340.

[8] Vijay Chidambaram, Yueh-Hsuan Chiang, and Bilge Mutlu. 2012. Designing persuasive robots: how robots might persuade people using vocal and nonverbal cues. In *Proceedings of the seventh annual ACM/IEEE international conference on Human-Robot Interaction*. ACM, 293–300.

[9] Robert B Cialdini and Noah J Goldstein. 2004. Social influence: Compliance and conformity. *Annu. Rev. Psychol.* 55 (2004), 591–621.

[10] Richard S Crutchfield. 1955. Conformity and character. *American Psychologist* 10, 5 (1955), 191.

[11] Morton Deutsch and Harold B Gerard. 1955. A study of normative and informational social influences upon individual judgment. *The journal of abnormal and social psychology* 51, 3 (1955), 629.

[12] Marlena R Fraune, Yusaku Nishiwaki, Selma Sabanović, Eliot R Smith, and Michio Okada. 2017. Threatening Flocks and Mindful Snowflakes: How Group Entitativity Affects Perceptions of Robots. In *Proceedings of the 2017 ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 205–213.

[13] Marlena R Fraune, Steven Sherrin, Selma Sabanović, and Eliot R Smith. 2015. Rabble of robots effects: Number and type of robots modulates attitudes, emotions, and stereotypes. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 109–116.

[14] Jaap Ham, René Bokhorst, Raymond Cuijpers, David van der Pol, and John-John Cabibihan. 2011. Making robots persuasive: the influence of combining persuasive strategies (gazing and gestures) by a storytelling robot on its persuasive power. In *International conference on social robotics*. Springer, 71–83.

[15] Jaap Ham and Cees JH Midden. 2014. A persuasive robot to stimulate energy conservation: the influence of positive and negative social feedback and task similarity on energy-consumption behavior. *International Journal of Social Robotics* 6, 2 (2014), 163–171.

[16] Bert H Hodges. 2014. Rethinking conformity and imitation: Divergence, convergence, and social understanding. *Frontiers in psychology* 5 (2014), 726.

[17] Bert H Hodges and Anne L Geyer. 2006. A nonconformist account of the Asch experiments: Values, pragmatics, and moral dilemmas. *Personality and Social Psychology Review* 10, 1 (2006), 2–19.

[18] Hideki Kozima, Marek P Michalowski, and Cocoro Nakagawa. 2009. Keepon. *International Journal of Social Robotics* 1, 1 (2009), 3–18.

[19] Iolanda Leite, Marissa McCoy, Monika Lohani, Daniel Ullman, Nicole Salomons, Charlene Stokes, Susan Rivers, and Brian Scassellati. 2015. Emotional storytelling in the classroom: Individual versus group interaction between children and robots. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 75–82.

[20] Iolanda Leite, Marissa McCoy, Daniel Ullman, Nicole Salomons, and Brian Scassellati. 2015. Comparing models of disengagement in individual and group interactions. In *Proceedings of the Tenth Annual ACM/IEEE International Conference on Human-Robot Interaction*. ACM, 99–105.

[21] Rosemarijn Looije, Mark A Neerincx, and Fokie Cnossen. 2010. Persuasive robotic assistant for health self-management of older adults: Design and evaluation of social behaviors. *International Journal of Human-Computer Studies* 68, 6 (2010), 386–397.

[22] Saul McLeod. 2008. Asch Experiment. (2008). https://www.simplypsychology.org/asch-conformity.html

[23] Kayako Nakagawa, Masahiro Shiomi, Kazuhiko Shinozawa, Reo Matsumura, Hiroshi Ishiguro, and Norihiro Hagita. 2013. Effect of robot's whispering behavior on people's motivation. *International Journal of Social Robotics* 5, 1 (2013), 5–16.

[24] Ahsan Nawroj, Mariya Toneva, Henny Admoni, and Brian Scassellati. An exploration of social grouping in robots: Effects of behavioral mimicry, appearance, and eye gaze. In *Proceedings of the Cognitive Science Society*, Vol. 36. 1060 – 1065.

[25] Masahiro Shiomi and Norihiro Hagita. 2016. Do Synchronized Multiple Robots Exert Peer Pressure?. In *Proceedings of the Fourth International Conference on Human Agent Interaction*. ACM, 27–33.

[26] Masahiro Shiomi, Kazuhiko Shinozawa, Yoshifumi Nakagawa, Takahiro Miyashita, Toshio Sakamoto, Toshimitsu Terakubo, Hiroshi Ishiguro, and Norihiro Hagita. 2013. Recommendation effects of a social robot for advertisement-use context in a shopping mall. *International Journal of Social Robotics* 5, 2 (2013), 251–262.

[27] Mikey Siegel, Cynthia Breazeal, and Michael I Norton. 2009. Persuasive robotics: The influence of robot gender on human behavior. In *IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2563–2568.