Under Pressure: Examining Social Conformity with Computer and Robot Groups

Nicholas Hertz & Eva Wiese

George Mason University

Manuscript type: Research Article

Word count: 4489

Acknowledgments: We thank Katherine Schaedel, Fariha Khan, and Asma Khan for their help with data collection. We also thank Ali Momen and Dr. Daniel Roberts for their assistance and insight.

ABSTRACT

Objective: To investigate whether non-human agents, such as computers or social robots, produce a social conformity effect within human operators and to what extent potential conformist behavior varies as a function of the human-likeness of the group members and the type of task that had to be performed. **Background:** People conform due to normative and/or informational motivations in human-human interactions, and conformist behavior is modulated by factors related to the individual, as well as factors associated with the group, context and culture. Studies have yet to examine whether non-human agents also induce social conformity. **Method:** Participants were assigned to a computer, robot, or human group and completed both a social and analytical task with the respective group. **Results:** Conformity measures (percentage of times participants answered in line with agents on critical trials) subjected to a 3 x 2 mixed ANOVA showed significantly higher conformity rates for the analytical versus the social task, as well as a modulation of conformity depending of the perceived agent-task fit. **Conclusion:** Findings indicate that non-human agents were able to exert a general conformity effect and that informational influence associated with the group's expertise for a given task had a stronger impact on conformity than normative motivations associated with its human-likeness. **Application:** Results suggest that users may react differently to suggestions of non-human versus human agent groups with the potential of under-reliance on social tasks.

Keywords: Compliance and Reliance, Trust in Automation, Human-Automation Interaction, Decision-Making, Social Psychology

Précis: Study examining whether groups of non-humans, computers and robot agents, are able to produce conformity effects in human interaction partners. Results indicate that conformity with each agent depends on the type of task being completed as well as the agent's perceived expertise in that task.

44 INTRODUCTION

45 Humans are social creatures and tend to organize themselves accordingly. From small groups
46 such as families or friend cliques to larger groups such as nations or societies, people are con-
47 stantly exposed to the influence of others. For almost a century, psychologists have studied the
48 effects of social influence on conformity behaviors of individuals. While Moore (1921) and She-
49 rif (1935) were the first to empirically explore the effects of social influence, Asch (1951) was
50 the first to specifically demonstrate social conformity in groups. Asch generated a disagreement
51 between an individual and a group on a perceptual task (i.e., line length judgments), and meas-
52 ured whether a consistently, but incorrectly answering majority had the power to bias individual
53 decision making towards prevailing group opinions. The group was composed of six confeder-
54 ates who were instructed to answer in a predefined pattern, placed in the same room as the partic-
55 ipant, and had to verbalize their answers in a serial pattern (i.e., one confederate at a time). It was
56 shown that the presence of a unanimous majority significantly biased the judgment of individuals
57 towards an incorrect response, with conformity rates as high as 33% (Asch, 1956). This is most
58 striking when considering that participants in the control condition (i.e., performing the task
59 alone) were extremely accurate at performing the task (< 1% error rate). Follow-up studies found
60 that conformist behavior was not only demonstrated for perceptual tasks like the line judgment
61 task, but also for subjective evaluative judgments (Alquist, Ainsworth & Baumeister, 2013; Bar-
62 on, Vandello & Brunsman, 1996; Cialdini & Goldstein, 2003) and memory tasks (Deuker et al,
63 2013; Edelson, Sharot, Dolan, & Dudai, 2011; Wright, Busnello, Buratto, & Stein, 2012). Con-
64 formist behavior was also observed when group members were only virtually present (i.e., on a
65 computer screen; Rosander & Eriksson, 2012), a set-up that was adopted in most modern con-
66 formity studies due to its simplicity (Schnuerch & Gibbons, 2014; Truatmann-Lengsfeld & Her-

mann, 2013). While social conformity has been consistently shown with human groups, it is unclear whether groups of non-human agents, like virtual avatars or social robots, would also have the potential to induce conformity in humans. However, this is an important question since interactions with non-human agents are becoming more required in everyday life. The goal of this paper is to examine a) whether humans conform to groups of non-human agents in a similar fashion as they conform to groups of other humans and b) to what extent social conformity depends on the human-likeness of the agent (i.e., computer vs. robot vs. human) and the nature of the task (i.e., social vs. analytical).

In human-human interactions, people conform to group beliefs due to *normative* and/or *informational* influences (Deutsch & Gerard, 1955; Schnuerch & Gibbons, 2014). Normative influence states that although people may not actually adopt the opinion of the group, they conform with others to fit in and to avoid social punishment that comes with standing out from the group (Levine, 1989; Schnuerch & Gibbons, 2014). Informational influence states that people may truly alter their beliefs to align with the group judgment and conform because they believe that the group knows more than them and that answering in line with the group is more likely to be correct (Cohen & Golden, 1972; Schnuerch & Gibbons, 2014). Both normative and informational explanations for social conformity are supported by neurocognitive research. For instance, public conformity is positively correlated with neural measures of anxiety (i.e., negative affective states associated with violations of group norms; Berns, Capra, Moore, & Noussair, 2010), conflict detection (i.e., detection of incongruences between self and group opinion; Chen, Wu, Tong, Guan, & Zhou, 2012), and reward processing (i.e., expectation of punishment if group opinion is not adopted; Klucharev, Hytönen, Rijpkema, Smidts, & Fernández, 2009). On the other hand, private conformity is associated with changes in memory (i.e., social influence causing

90 increased limbic activity that correlates with long-lasting memory changes; Edelson et al., 2011)
91 and object value representations (i.e., true attitude changes rather than capitulation to group opin-
92 ion; Mason, Dyer, & Norton, 2009; Zaki, Schirmer, & Mitchell, 2011). Conforming to a group is
93 also associated with detections of cognitive inconsistency between self and group judgment
94 (Klucharev et al., 2009), and a disparity between own preferences and what is considered to be
95 socially desirable (Izuma & Adolphs, 2013; Schnuerch & Gibbons, 2014).

96 Aside from normative and informational influences, a number of contextual and disposi-
97 tional factors have been shown to influence social conformity. An important factor influencing
98 the presence of social conformist behavior is *group size*, with initial increase in social conformity
99 as group size increases, followed by a diminished effect once group size has a number greater
100 than about three (Asch, 1956; see Bond & Smith, 1996; for a meta-analysis). It has also been
101 shown that *similarity* with a group (i.e. ingroup status) consistently increased conformity com-
102 pared to an outgroup status (Abrams, Wetherell, Cochrane, Hogg, & Turner, 1990), though
103 McKelvey & Kerr (1988) demonstrated that people were significantly less likely to conform in
104 the presence of friends than strangers. In addition, factors like *anonymity* of the response scheme
105 (i.e., lower conformity for responses anonymous responses; Allen, 1965), *severity of deviation*
106 from the group opinion (i.e., higher conformity the greater the deviation; Tanford & Penrod,
107 1984), or *participant gender* (i.e., higher conformity for females than males; Cooper, 1979) also
108 affected social conformity. Notably, however, the gender effect on conformity seems to be
109 strongly influenced by historical and cultural factors and has not been replicated in more recent
110 studies (Rosander & Eriksson, 2012). People growing up in *collectivist cultures* (i.e., social be-
111 havior is largely determined by group goals) were also more likely to conform to group beliefs

than individuals from individualist cultures (i.e., social behavior is largely determined by personal goals; Schwartz, 1992; see Bond & Smith, 1996; for a meta-analysis).

Although previous research has shown that factors associated with the individual (e.g., gender), the group members (e.g., similarity), and the context (e.g., anonymity of the response setting) modulate social conformity, it has not been fully determined yet if non-human agents can also induce social conformity, and if so, to what extent conformist behavior is modulated by the human-likeness of the group and the type of task that has to be decided on. The two studies that were conducted on this topic so far were unable to show social conformist behavior when the group consisted of robotic agents (Beckner et al., 2016; Brandsetter et al., 2014), but have not looked at the role agent expertise plays for social conformity (i.e., agent-task fit). They have also not distinguished between normative (i.e., conforming to avoid rejection) and informational influence (i.e., conforming based on group expertise), although one would expect human-likeness to differentially affect these variables. The potential of non-human agents to induce social pressure and exert normative influence should be reduced compared to human groups owing to a reduced capacity to provide social approval in case of conformist behavior, whereas their potential to exert informational influence should be modulated depending on whether they are considered to be competent for a given task (e.g., competent/incompetent on analytical/social tasks), making it likely that non-human agents could induce social conformity at least in technical domains. The assumption that non-human groups can trigger conformist behavior based on expertise (i.e., informational) rather than social approval (i.e., normative) would be in line with observations by Gray, Gray & Wegner (2007), who showed that although robots are ascribed the ability to plan (i.e., agency), they are denied the ability to feel (i.e., experience), which makes them powerful decision-makers, but weak social approvers. To test these assumptions, we ask participants in the

current experiment to perform a social (i.e., judging emotional states) and an analytical task (i.e., standard arithmetic operators) with groups of agents varying in their human-likeness from computer to social robot to human. We expect a general increase in conformity as the human-likeness of the group increases due to a higher potential of human agents to exert normative influence. In addition, we expect the effect of human-likeness on conformity to be modulated by task type due to informational influence, that is: human groups should be deemed competent for both the social and the analytical task, while machine agents should only be considered to be competent to perform the analytical task.

METHODS & MATERIALS

Participants

90 undergraduate students (59 women; $M_{\text{age}} = 19.85$ years, range: 18-39) at George Mason University participated in the experiment. Participants were recruited via SONA (i.e., online study pool recruitment platform), and compensated for their participation through course credit. Participants were randomly assigned to one of the three agent conditions at the beginning of the experiment: 30 participants took part in the *Computer* condition (23 women, $M_{\text{age}} = 19.4$), 29 participants in the *Robot* condition (18 women, $M_{\text{age}} = 21.0$), and 31 participants in the *Human* condition (18 women, $M_{\text{age}} = 19.2$). Task type was manipulated within-participants, that is: all participants performed both the analytical task and the social task with their assigned agent.

Apparatus

The stimuli were presented on a 21.5-inch iMac with the refresh rate of 60 Hz. Participants were seated approximately 57 cm from the monitor and experimenters ensured the participant was centered in respect to the monitor. The experiment was programmed in PsychoPy.

Task

Participants performed two tasks during the experiment: a social task and an analytical task. As social task, we used the *Reading the Mind in the Eyes* task (Baron-Cohen, Wheelwright, Hill, Raste, & Plumb, 2001), which assesses the ability to recognize complex emotions in still images of the eye region of human faces. The task consists of 36 images of the eye-region of different actors and actresses, and participants are asked to choose which of four presented words best described what the person in the photograph was feeling. The analytical task required participants to solve standard addition and subtraction operations with a series of dot stimuli presented on the screen (adapted from Smith, Allaham, & Wiese, 2016). Social and analytical task were presented simultaneously throughout the experiment, with the images of the eye regions of different faces (i.e., social task) being shown in the center of the screen, and the series of dots being presented above and below the image (i.e., analytical task).

To complete the social task, participants had to determine what emotion or mental state the person in the image was depicting (e.g. “What emotion was expressed?”). To complete the analytical task, participants had to complete an addition or subtraction operation using the white dots surrounding the social stimulus (e.g. “Top Dots + Bottom Dots =”). In both tasks, participants would respond by selecting one of four multiple-choice options located at each corner of the screen.

Agents

Participants were told that they could get advice from one of three types of agent groups while working on the two tasks: computers, robots, or humans (which were introduced via short cover stories; see [Supplementary Materials](#)). Participants were told that the groups would consist

of four agents that were randomly selected from a database of agents who already completed the experiment in September 2016, and that the actions they performed and answers they gave back then on a trial-by-trial basis were video recorded and would be replayed for the participants as an aid for their own decision-making on respective trials.

In the *human* condition, participants were additionally told that the former participants were students at GMU and that the video would show their hand and arm movements, a screen capture of their answer selection, and the answer they ultimately picked on each trial. Four individual human hands were used throughout the experiment to cause the impression that the choices of different individuals were shown (**Figure 1**).

Figure 1. *Agent Representations:* Still-shot examples of the computer, robot, and human agent videos used in the experiment. Computer and robot videos were created entirely in Adobe Photoshop and After Effects. The human videos were recorded from the same angle as the deceptive cameras location in the experiment.

The videos were created by recording four lab members while performing the task: two males and two females. They were instructed how to answer in order to obtain the desired response pattern (see Asch, 1956). To make the manipulation more believable, participants were told that their hand and mouse movements would be recorded and then displayed to participants who would take part in the experiment later in the semester. For that purpose, we installed a

camera above the location where the mouse was placed to cause the impression that the video was recorded from the same angle as was shown in the videos they watched themselves; see

Figure 2.

Figure 2. *Laboratory set-up:* Image of the computer station at which participants performed the experiment. All critical measures were completed using the computer mouse. Participants were told their hand movements and screen would be recorded and displayed to future participants to make the manipulation more believable. For this purpose, a video camera was placed above the mouse and directed at the participant's hand. In reality, this camera was nonoperational and did not actually record participant data.

In the *robot* condition, the four agents were described as having been created using 3D-printing technology, and as being programmed to have similar cognitive abilities as computers. Similar to the human condition, the robot agents were represented by four different robotic arms controlling a computer mouse and ultimately selecting one out of the four responses. The robot hands were created by combining images of a biomimetic robotic hand and forearm (Todorov &

Xu, 2016) in Adobe Photoshop and selectively changing the color of the fingertips to represent four different robots (**Figure 1**). Videos of the answer selection process were created using key-frame animation in Adobe After Effects.

In the *computer* condition, the agents were described as being equipped with algorithms designed to learn and answer the cognitive tasks presented in the experiment, and their actions were depicted as scrolling computer code when selecting a final response. The computer agents were only distinguishable by their location on the participants screen when completing the tasks. The videos of the agents were created using a combination of Adobe Photoshop and After Effects.

Procedure

At the beginning of the experiment, participants were randomly assigned to one of the three agent conditions (i.e., human, robot or computer), gave informed consent, and completed a brief demographic questionnaire. Participants were then introduced to the agent group they were assigned to via cover stories, and received general instructions as to how to perform the social and analytical tasks: they were told that after the presentation of the stimulus a prompt for which task they were to complete would appear, together with four possible answers, presented in the four corners of the screen (either emotional states that can be inferred from the eye region of a human face or the solution to an addition or subtraction problem). Participants were also told that their actions and answers would be recorded and eventually replayed to participants in future experiments. The deceptive camera was placed above the participants mouse hand (which was non-functioning and did not actually record participants' actions). If participants had no further questions, the test phase began.

Participants were presented with 36 trials in total: 18 analytical trials (i.e., arithmetic operations involving dots) and 18 social trials (i.e., judging emotional states based on eye region). On each trial, the analytical and social task were presented simultaneously, that is: an image depicting the eye region of a face showing a certain emotion was presented in the center of the screen, with a number of white dots presented above and below the image. Participants had to pay attention to both the social and analytical stimuli, since it was only after the presentation of these stimuli that participants were told whether they had to solve the analytical or the social task (this manipulation was chosen to assure that both sets of stimuli were processed). The accuracy with which the agents answered the questions followed the pattern of the traditional conformity paradigm (Asch, 1956): within 36 trials, there were 12 neutral trials (i.e., agents answered unambiguously correctly; 6 per task type) and 24 critical trials (i.e., agents answered unambiguously incorrectly; 12 per task type), and the order in which correct and incorrect answers were given followed a certain pattern (i.e., Neutral, Neutral, Critical, Critical, Neutral, Critical, Critical, Critical, Critical, Critical; repeated). The order in which the analytical and social questions were presented was randomized across 18 trials, making sure there were equal numbers of neutral and critical trials for each task type. This pattern was then mirrored for the next 18 trials.

The sequence of events on a given trial was as follows: at the beginning of each trial the stimuli for the social task (i.e., image of eye region) and for the analytical task (i.e., number of dots) were simultaneously presented for 2.5 seconds. Afterwards, either the social (e.g., “What emotion was expressed?”) or the analytical (e.g., “Top dots + Bottom dots =”) task was presented together with the four possible answers (presented in the corners of the screen; see **Figure 3**).

Figure 3. Trial Sequence: At the beginning of each trial the social and analytical task were presented on screen for 2.5 seconds (analytical task presented in this example). Afterwards, participants observed four agents answer sequentially (human agent in the third position presented in the figure). After having observed all four agents select their answers, the participant was then instructed to choose their answer.

248

249 Before participants could select their response, four videos depicting each of the four agents per

250 group answering the respective question auto-played. Each video was approximately 3 seconds

251 long, and all videos were played sequentially. At the end of each video, the answer the respective

252 agent chose appeared on the screen, and participants were prompted to give their answers. All

253 agent answers remained on the screen until participants had selected their response to the ques-

254 tion, which prompted the next trial. The experiment took approximately 30 minutes and the Insti-

255 tutional Review Board (IRB) at George Mason University approved procedures before the start

256 of the study.

RESULTS

Conformity was determined as the percentage of times participants answered in line with the agents on critical trials (i.e., trial was answered unanimously incorrectly), which were then subjected to a 3 x 2 mixed ANOVA with *Agent* (Computer vs. Robot vs. Human) as a between-participants factor and *Task* (Social vs. Analytical) as a within-participants factor. The ANOVA showed a main effect for *Task* ($F(1, 85) = 25.282, p < .001, \eta^2 = .044$), with significantly higher conformity rates for the analytical versus the social task (46.51% vs. 35.56%). The main effect *Agent* was not significant ($F(2, 85) = 1.666, p = .195, \eta^2 = .032$), showing that conformity was comparable across agent types (Computer: 35.56%; Robot: 40.37%; Human: 47.18%). Most importantly, however, the *Agent* x *Task* interaction was significant ($F(2, 85) = 4.738, p = .001, \eta^2 = .0169$), showing that the type of task participants performed affected how strongly they trusted different agent types (see **Figure 4**).

Post-hoc univariate ANOVAs (one for the analytical and one for the social task) showed no main effect of *Agent* for the analytical task ($F(2, 85) = .315, p = .731, \eta^2 = .007$), but a significant main effect of *Agent* for the social task ($F(2, 85) = 4.169, p = .0187, \eta^2 = .0893$), with significantly higher conformity in the human condition than in the machine conditions (i.e., computer and robot conditions combined; $t(87) = 2.374, p = .0198$).

Figure 4. *Mean conformity measures by agent and task type.* Analyses revealed a main effect of task type, with stronger conformity during analytical than social trials. No main effect of agent type was found, indicating that human and non-human agents exerted comparable amounts of social pressure across task types (i.e., comparable normative influence). However, the interaction of agent type and task type was significant and post-hoc tests revealed that conformity with the machine agents on the social task was significantly reduced compared to the human agents, indicating a calibration of conformist behavior according to the agents' perceived task expertise (i.e., informational influence). Error bars represent standard errors and points are offset horizontally so that error bars are visible.

DISCUSSION

The goal of this study was to investigate whether non-human agents, such as computers and social robots, are able to exert enough social pressure to produce a conformity effect within human partners and whether the human-likeness of the group members and the type of task that had to be performed modulated this effect. We hypothesized that high levels of human-likeness

would lead to high conformity across the board (due to normative influences) and that task type would modulate the degree of conformity according to group competencies, with high conformity with human agents on both the social and the analytical task, high levels of conformity with robot and computer agents on the analytical task, but reduced levels of conformity on the social task (due to informational influences). The findings support these hypotheses: there was a significant interaction effect between agent type and task type on conformity measures, showing that participants adjusted their conformist behavior according to the group's competence for a given task. Specifically, participants conformed to comparable degrees with human, robot and computer groups during the analytical task (in line with the assumption that all agents are capable of performing relatively simple math operations), but conformed significantly more strongly with the group of humans on the social task than the group of robots or computers (in line with the assumption that machines have reduced abilities to perceive emotions; Gray et al., 2007). Interestingly, no main effect of agent type was found in the current experiment, showing that a group of human agents did not generally induce more social conformity than groups of non-human agents. Taken together, this pattern of findings indicate that informational influence associated with the group's expertise for a given task had a strong impact on social conformity, while normative motivations associated with social rejection by the group did not modulate this effect further.

The finding that participants conform more to human than non-human agents on social tasks is in line with observations by Gray and colleagues (2007) that humans are seen as being capable of rational thought and planning (i.e., agency) as well as feeling and identifying emotions (i.e., experience), while robot and computer agents are only ascribed agency, but not experience. Consequently, human agents would be able to exert informational influence on both ana-

lytical and social tasks, while robot and computer agents would be expected to do so only for analytical tasks (which is in line with the reported findings). Similarly, it was shown that people react more strongly to the inaccuracies of computers than humans (e.g., Dzindolet, Pierce, Beck, Dawe, & Anderson, 2001), show a faster decline in trust after computer errors (Wiegmann, Rich, & Zhang, 2001), and tend to view computer support systems as more objective and rational than humans (Dijkstra, Liebrand, & Timminga, 1998). Thus participants may have responded more sensitively to the machine inaccuracies in the current experiment (agent groups answered correctly 33% of the time), in particular on a task that does not necessarily fall within their core competencies, like the social task. Finding machines as being more objective and rational than humans would explain why conformity to machine agents is not significantly different from conformity to human agents, but differs for the social task (i.e., effect mainly due to differences in informational influence).

The results of this study appear to be in contrast with other studies that have examined social conformity with groups of non-human agents: While the current experiment shows significant conformity effects for all agents, Beckner and colleagues (2016) found that humanoid robots did not induce social conformity effects on a verbal task (i.e., production of English past tense). Likewise, Brandsetter and colleagues (2014) showed that robot groups were not able to exert social pressure on perceptual (i.e., line judgments) and verbal tasks (i.e., production of English past tense). However, in light of the observation that conformity with non-human agents was task-dependent, it is possible that the tasks that were used in previous studies did not provide a good enough agent-task fit when interacting with robot groups thereby failing to induce informational influence. Another possibility is that static group representations (i.e., agents represented via still images), as used in previous experiments, are in general less potent at inducing normative influ-

ence because they do not allow social rejection by group members. This interpretation would be in line with results from the social facilitation literature, where social robots (Riether, Hegel, Wrede, & Horstmann, 2012; Woods, Dautenham, & Kaouri, 2005) and avatars (Hoyt, Blascovich, & Swinth, 2003; Park & Catrambone, 2007) were able to induce social facilitation but only if they displayed a certain level of embodiment (i.e., static images having less influence than moving images having less influence than agents physically present in the room).

There are some limitations associated with the current experiment: First, the assumption that human agents are associated with high analytical and social skills, while machine agents are only associated with high analytical skills is exclusively based on theoretical assumptions (Andrist, Spannan, & Mutlu, 2013; Chidambaram, Chiang, & Mutlu, 2012; Gray et al., 2007). That is, actual beliefs about the social and analytical capabilities of the different agent types or general attitudes towards robotic agents were not assessed. The reason why we did not obtain these measures is that asking participants about the different agents' skill sets could potentially lead to biased data: measuring participants' beliefs before the task, might have made their pre-existing beliefs unnecessarily explicit or revealed too much information about the experimental manipulation; while assessing beliefs after the interaction would have been biased by the interaction with the agents and not accurately reflect pre-existing assumptions. Future research should address this limitation by explicitly manipulating beliefs regarding the agents' expertise via instructions and by assessing a-priori beliefs about the social and analytical abilities of computers outside the experimental setting (similar to Gray et al., 2007). Second, it cannot ultimately be determined to what extent the interaction between agent and task type was due to normative or informational influence. On the one hand, it is possible that human groups induce conformity both due to normative and informational motivations, while machine groups only exert informational

influence. If that were to be the case, machines would exert higher informational influence on analytical tasks than human agents, but since they exert weaker normative influence, the overall conformity with human versus non-human agents during analytical tasks would not differ. For social tasks, human groups would exert both informational and normative influence, while machine groups could only exert informational influence (if at all), resulting in lower conformity with machine groups (as is the case in the current experiment). While this explanation is possible, it is unlikely because even the least human-like group, namely the computer agents, triggered conformity rates of about 25% (i.e., only 8% lower than in Asch, 1956) on a task that they were not stereotypically competent in (i.e., social task), which indicates that they must have exerted normative influence at least to some extent. Alternatively, it is possible that participants found the social task difficult enough (i.e., perception of own abilities is low) and/or the computers capable enough for this task that even their limited expertise was considered helpful enough to induce informational influence, which caused an observational conformity rate despite the potential of machine agents lack of ability to induce normative influence. Thus, it cannot ultimately be determined whether the observed interaction effect is due to variations in informational influence, assuming comparable normative influence for all agent types (i.e., human competent for both tasks, machines only for the analytical task) or due to an interplay between informational and normative influences (i.e., human/machines competent for social/analytical tasks, but normative influence due to rejection is only feared in the human condition). To elucidate the impact of informational versus normative influences on conformity with mechanistic agents, further research is needed that allows for manipulating normative influences directly via public versus private response schemes.

The results of the current experiment have implications for human-machine-teams, design of decision aid systems, and many other systems that involve human-machine interaction. For example, human operators have to work together with groups of automated agents in military settings (i.e., loyal wingman scenarios) or in modern workplaces (i.e., for companies like Amazon), may encounter situations in which informational and normative influences bias their decisions in physically or financially dangerous situations. Designers of these systems should be aware that operators will react differently to human versus non-human group influences depending on the task and consider calibrating the human-likeness of the system based on the types of tasks the operator will most likely have to execute.

Key points

- Investigated whether non-human agents, such as computers and social robots, produce a conformity effect within human partners
- People conform due to either normative or informational influence and previous research has shown that factors associated with both the individual, the group, contextual, and societal factors modulate social conformity.
- Participants were assigned to a computer, robot, or human group condition and completed both a social and analytical task with those agents.
- Found significantly higher conformity rates for the analytical versus the social task, but that is modulated by an interaction between type of task and agent.
- Participants conformed to comparable degrees with human, robot and computer groups during the analytical task, but conformed significantly more strongly on the social task as the group's human-likeness increased.

REFERENCES

- Abrams, D., Wetherell, M., Cochrane, S., Hogg, M. A., & Turner, J. C. (1990). Knowing what to think by knowing who you are: Self-categorisation and the nature of norm formation, conformity and group polarisation. *British Journal of Social Psychology*, 29, 97-119.
- Allen, V. L. (1965). Situational factors in conformity. In L. Berkowitz (Ed.), *Advances in experimental social psychology* (Vol. 2, pp. 133-175). New York: Academic Press.
- Alquist, J. L., Ainsworth, S. E., & Baumeister, R. F. (2013). Determined to conform: Disbelief in free will increases conformity. *Journal of Experimental Social Psychology*, 49, 80-86.
- Andrist, S., Spannan, E., & Mutlu, B. (2013, March). Rhetorical robots: making robots more effective speakers using linguistic cues of expertise. In *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction* (pp. 341-348). IEEE Press.
- Asch, S. (1951). Effects of group pressure upon the modification and distortion of judgments. In H. Guetzkow (Ed.), *Groups, leadership and men* (pp. 177-190). Oxford, UK: Carnegie Press.
- Asch, S. (1956). Studies of independence and conformity: I. a minority of one against a unanimous majority. *Psychological Monographs: General and Applied*, 70(9), 1-70.
- Baron, R. S., Vandello, J. A., & Brunsman, B. (1996). The forgotten variable in conformity research: Impact of task importance on social influence. *Journal of Personality and Social Psychology*, 71(5), 915-927.
- Baron-Cohen, S., Wheelwright, S., Hill, J., Raste, Y., & Plumb, I. (2001). The "Reading the Mind in the Eyes" test revised version: A study with normal adults, and adults with Asperger syndrome or high-functioning autism. *Journal of Child Psychology and Psychiatry*, 42(2), 241-251.

- 417 Beckner, C., Rácz, P., Hay, J., Brandsetter, J., & Bartneck, C. (2016). Participants conform to
418 humans but not to humanoid robots in an English past tense formation task. *Journal of*
419 *Language and Social Psychology*, 35(2), 158-179.
- 420 Berns, G. S., Capra, C. M., Moore, S., & Noussair, C. (2010). Neural mechanisms of the influ-
421 ence of popularity on adolescent ratings of music. *NeuroImage*, 49, 2687–2696.
- 422 Bond, R., & Smith, P. B. (1996). Culture and conformity: A meta-analysis of studies using
423 Asch's (1952b, 1956) line judgment task. *Psychological bulletin*, 119(1), 111.
- 424 Brandstetter, J., Rácz, P., Beckner, C., Sandoval, E. B., Hay, J., & Bartneck, C. (2014, Septem-
425 ber). A peer pressure experiment: Recreation of the Asch conformity experiment with ro-
426 bots. In *Intelligent Robots and Systems (IROS 2014), 2014 IEEE/RSJ International Con-*
427 *ference on* (pp. 1335-1340). IEEE.
- 428 Chen, J., Wu, Y., Tong, G., Guan, X., & Zhou, X. (2012). ERP correlates of social conformity in
429 a line judgment task. *BMC neuroscience*, 13(1), 43.
- 430 Chidambaram, V., Chiang, Y. H., & Mutlu, B. (2012, March). Designing persuasive robots: How
431 robots might persuade people using vocal and nonverbal cues. In *Proceedings of the sev-*
432 *enth annual ACM/IEEE international conference on Human-Robot Interaction* (pp. 293-
433 300). ACM.
- 434 Cialdini, R. B., & Goldstein, N. J. (2003). Social influence: Compliance and conformity. *Annual*
435 *Review of Psychology*, 55, 591-621.
- 436 Cohen, J. B. & Golden, E. (1972). Informational social influence and product evaluation. *Journal*
437 *of Applied Psychology*, 56, 54-59.

- 438 Cooper, H. M. (1979). Statistically combining independent studies: A meta-analysis of sex dif-
439 ferences in conformity research. *Journal of Personality and Social Psychology*, 37, 131-
440 146.
- 441 Deuker, L., Müller, A. R., Montag, C., Markett, S., Reuter, M., Fell, J., Trautner, P., & Axmach-
442 er, N. (2013). Playing nice: A multi-methodological study on the effects of social con-
443 formity on memory. *Frontiers in Human Neuroscience*, 7, 79.10.3389/fnhum.2013.00079
- 444 Deutsch, M., & Gerrard, H. B. (1955). A study of normative and informational social influences
445 upon individual judgment. *Journal of Abnormal and Social Psychology*, 51, 629-636.
- 446 Dijkstra, J. J., Liebrand, W. B. G., & Timminga, E. (1998). Persuasiveness of expert systems.
447 *Behavior & Information Technology*, 17(3), 155-163.
- 448 Dzindolet, M. T., Pierce, L. G., Beck, H. P., Dawe, L. A., & Anderson, B. W. (2001). Predicting
449 misuse and disuse of combat identification systems. *Military Psychology*, 13(3), 147-164.
- 450 Edelson, M., Sharot, T., Dolan, R. J., & Dudai, Y. (2011). Following the crowd: Brain substrates
451 of long-term memory conformity. *Science*, 333, 108–111.
- 452 Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*,
453 315(5812), 619.
- 454 Hoyt, C. L., Blascovich, J., & Swinth, K. R. (2003). Social inhibition in immersive virtual envi-
455 ronments. *Presence: Teleoperators & Virtual Environments*, 12(2), 183-195.
- 456 Izuma, K., & Adolphs, R. (2013). Social manipulation of preference in the human brain. *Neuron*,
457 78, 563–573.
- 458 Klucharev, V., Hytönen, K., Rijpkema, M., Smidts, A., & Fernández, G. (2009). Reinforcement
459 learning signal predicts social conformity. *Neuron*, 61, 140–151.

- 460 Levine, J. M. (1989). Reaction to opinion deviance in small groups. In P. B. Paulus (Ed.), *Psy-*
461 *chology of group influence* (2nd ed., pp. 187-231). Hillsdale, NJ: Erlbaum.
- 462 Mason, M. F., Dyer, R., & Norton, M. I. (2009). Neural mechanisms of social influence. *Organi-*
463 *zational Behavior and Human Decision Processes*, 110, 152–159.
- 464 McKelvey, W., & Kerr, N. H. (1988). Differences in conformity among friends and strangers.
465 *Psychological Reports*, 62(3), 759-762.
- 466 Moore, H. T. (1921). The comparative influence of majority and expert opinion. *American Jour-*
467 *nal of Psychology*, 32, 1620.
- 468 Park, S., & Catrambone, R. (2007). Social facilitation effects of virtual humans. *Human Fac-*
469 *tors*, 49(6), 1054-1060.
- 470 Riether, N., Hegel, F., Wrede, B., & Horstmann, G. (2012, March). Social facilitation with social
471 robots?. In *Proceedings of the seventh annual ACM/IEEE international conference on*
472 *Human-Robot Interaction* (pp. 41-48). ACM.
- 473 Rosander, M., & Eriksson, O. (2012). Conformity on the internet – the role of task difficulty and
474 gender differences. *Computers in Human Behavior*, 28, 1587–1595.
- 475 Schnuerch, R. & Gibbons, H. (2014). A review of neurocognitive mechanisms of social con-
476 formity. *Social Psychology*, 45(6), 466-478.
- 477 Schwartz, S. H. (1992). Cultural dimensions of values: Towards an understanding of national
478 differences. In U. Kim, H. C. Triandis, C. Kagitcibasi, S. C. Choi, & G. Yoon (Eds.), *In-*
479 *dividualism and collectivism: Theory, method and applications* (pp. 85-119). Thousand
480 Oaks, CA: Sage.
- 481 Sherif, M. (1936). *The psychology of social norms*. New York, NY: Harper.

- 482 Smith, M. A., Allaham, M. M., & Wiese, E. (2016, September). Trust in automated agents is
483 modulated by the combined influence of agent and task type. In *Proceedings of the Hu-*
484 *man Factors and Ergonomics Society Annual Meeting* (Vol. 60, No. 1, pp. 206-210).
485 Sage CA: Los Angeles, CA: SAGE Publications.
- 486 Tanford, S. & Penrod, S. (1984). Social influence model: A formal integration of research on
487 majority and minority influence processes. *Psychological Bulletin*, 95, 189-225.
- 488 Todorov, E., & Xu, Z. (2016). Design of a highly biomimetic anthropomorphic robotic hand to-
489 wards artificial limb regeneration. *2016 IEEE International Conference on Robotics and*
490 *Automation (ICRA)*, 3485-3492.
- 491 Trautmann-Lengsfeld, S. A., & Herrmann, C. S. (2013). EEG reveals an early influence of social
492 conformity on visual processing in group pressure situations. *Social Neuroscience*, 8, 75–
493 89.
- 494 Wiegmann, D. A., Rich, A., & Zhang, H. (2001). The effects of aid reliability on users' trust and
495 reliance. *Theoretical Issues in Ergonomics Science*, 2(4), 352-367.
- 496 Woods, S., Dautenhahn, K., & Kaouri, C. (2005). Is someone watching me? - consideration of
497 social facilitation effects in human-robot interaction experiments. *CIRA 2005, IEEE In-*
498 *ternational Symposium on Computational Intelligence in Robotics and Automation*. 53–
499 60.
- 500 Wright, D. B., Busnello, R. H. D., Buratto, L. G., & Stein, L. M. (2012). Are valence and social
501 avoidance associated with the memory conformity effect? *Acta Psychologica*, 141, 78–
502 85.
- 503 Zaki, J., Schirmer, J., & Mitchell, J. P. (2011). Social influence modulates the neural computa-
504 tion of value. *Psychological Science*, 22, 894–900.

505 **Biographies:**

506 Nicholas Hertz is a PhD student in the Human Factors and Applied Cognition Program at George
507 Mason University, where he received his M.A. in 2014.

508

509 Eva Wiese is an Assistant Professor in the Human Factors and Applied Cognition Program at
510 George Mason University. She received her PhD in Neuroscience from the Ludwig-Maximilian
511 University of Munich, Germany in 2013.