**Implicit mind perception alters vigilance performance
due to cognitive conflict processing**

Eva Wiese^a, Arielle Mandell^{a*}, Tyler Shaw^a & Melissa Smith^a

^aPsychology Department, George Mason University
4400 University Drive, Fairfax, VA 22030, USA

*Corresponding author. Email: amandel3@gmu.edu
Address: 4400 University Dr. MS3F5, Fairfax, VA 22030, USA

ABSTRACT

Knowing the internal states of others is essential to predicting behavior in social interactions and requires that the general characteristic of 'having a mind' is granted to our interaction partners. Mind perception is a highly automatic process and can potentially cause a cognitive conflict when interacting with agents whose mind status is ambiguous, such as artificial agents. We investigate whether mind perception negatively impacts performance on tasks involving artificial agents due to cognitive conflict processing caused by a potentially increased difficulty to categorize them as human versus non-human. Experiment 1 shows that an ambiguous humanoid stimulus negatively impacts performance on a vigilance task that is known to be sensitive to the drainage of cognitive resources. This negative effect on performance vanishes when participants are pre-exposed to the stimulus before the vigilance task (Experiment 2 and 3). The effect of pre-exposure on performance recovery is independent of whether participants explicitly resolve the cognitive conflict by answering mind-related questions (Experiment 2) or implicitly by judging the stimuli on a set of physical features (Experiment 3). Together, the findings suggest that mind perception is so automatic that it cannot be suppressed even if it has negative effects on cognitive performance.

Significance Statement

This study suggests that the automatic process of assessing whether or not an entity has a mind can have negative effects on sustained attention when interacting with entities that are not easily categorized as human or non-human, such as very human-like robots. This negative effect disappears when people are given the opportunity to assess the mental or physical attributes of the entity before interacting with it.

Keywords

Mind perception, vigilance, performance, cognitive conflict, human-robot interaction

INTRODUCTION

In social interactions, we attribute internal states to others and activate brain areas responsible for social cognition in order to understand and predict the behavior of others (Frith & Frith, 2006). The same brain networks used to understand our own behavior are used to simulate the behavior of others in order to derive information about their motives, emotions and intentions (e.g., Ciaramidaro, Adenzato, Enrici, Erk, Pia, et al., 2007). This basic process of making inferences about internal states of others is referred to as mentalizing, and presupposes that the entity whose behavior we want to understand is perceived to have a mind (Premack & Woodruff, 1978; Baron-Cohen, Leslie, & Frith, 1985; Dennett, 2003)¹. *Mind perception* is not exclusive to agents that actually have a mind, but also occurs for agents that by definition do not have minds, such as robots and other artificial agents (Gray, Gray & Wegner, 2007). Imbuing nonhuman entities with human features is the essence of anthropomorphism, and is the default mechanism when novices try to explain the behavior of artificial agents but lack system-specific knowledge (Epley, Waytz, & Cacioppo, 2007; Rosset, 2008). Due to anthropomorphism, we believe that pets play tricks on us when they hide our personal items, imagine spiritual deities as having a human physique, and expect artificial agents to oppress mankind once they have the ability to outperform us (Eyssel & Reich, 2013). In this paper, we ask the fundamental question of whether the ten-

¹ Please note that mind perception is distinct from mentalizing, but related to the “problem of other minds”: *Mind perception* is the process by which an agent is deemed as generally capable of having internal states based on its appearance or behavior (Gray et al., 2007), while *mentalizing* is the process of inferring particular internal states of agents that have been labeled as ‘having a mind’ (Baron-Cohen, 1985). This distinction is similar to the *problem of other minds*, which is concerned with the question of what entitles us to our belief that others have inner lives rather than with the question of what specifically is happening in those inner lives (Wittgenstein, 1953).

dency to perceive mind in nonhuman agents has a negative effect on cognitive performance on tasks involving agents with ambiguous mind status due to an increased difficulty to categorize them as human versus nonhuman.

When we perceive minds in other entities, we adopt an intentional stance and interpret their behavior as goal-directed and intentional (Dennett, 2003). Agents believed to be devoid of intentional behavior are deemed incapable of experiencing complex emotional or motivational states and activate social brain areas to a lesser extent than intentional agents (e.g., Leyens et al., 2000; Sanfey, Rilling, Aronson, Nystrom, & Cohen, 2003; Harris & Fiske, 2006; Demoulin et al., 2009). These brain areas include anterior areas such as the dorsolateral, medial prefrontal and anterior cingulate cortex (DLPFC, MPFC, ACC) involved in mentalizing and empathizing, as well as posterior areas such as the temporo-parietal junction (TPJ), fusiform face area (FFA) and amygdala responsible for mentalizing, emotional processing and face perception (van Overwalle, 2009; for a review). As a consequence, people are less empathetic, display fewer humanitarian acts, and show less helping behaviors towards agents not perceived to have a mind (Haslam, 2006; Hein, Silani, Preuschoff, Batson, & Singer, 2010; Cikara, Bruneau, & Saxe, 2011; Harris & Fiske, 2011; Gutsell & Inzlicht, 2012), while showing more prosocial behaviors such as increased generosity or decreased cheating and reporting deeper social connections when interacting with agents perceived to have a mind (e.g., Haley & Fessler, 2005; Bering & Johnson, 2005; Shariff & Norenzayan, 2007; Gray et al., 2007; Epley et al., 2008 a, b; Gray, Young & Waytz, 2012).

Mind perception is a highly automatic process that activates social brain areas in an implicit or bottom-up fashion (e.g., Looser & Wheatley, 2010; Gao, McCarthy & Scholl, 2010; Wheatley, Weinberg, Looser, Moran, & Hajcak, 2011; Schein & Gray, 2015) if an entity pos-

sesses physical features of “humanness”, such as prototypical facial elements and configurations (e.g., eye-nose-mouth configuration; Maurer, Le Grand, & Mondloch, 2002; Gao et al., 2010; Looser & Wheatley, 2010; Balas & Tonsager, 2014; Schein & Gray, 2015; Deska, Almaraz & Hugenberg, 2016), or biological motion patterns (e.g., Castelli, Happé, Frith, & Frith, 2000). As a result, intentional agents can be differentiated from mechanistic agents within a few hundred milliseconds (Wheatley, Weinberg, Looser, Moran, & Hajcak, 2011; Looser, Guntupalli, & Wheatley, 2013), and even passively viewing stimuli triggering mind perception is sufficient to activate social brain networks (Wagner, Kelley, & Heatherton, 2011), even if their mind status is task-irrelevant (Wiese et al., 2014; Wykowska et al., 2014; Martini et al., 2015). “Mind” can also be ascribed to others in an explicit or top-down fashion when the presence of a social entity is needed or when an entity has become so important to an individual that a “machine” status would not be sufficient: nonhuman entities are more likely treated as “human” when individuals are in increased need of social contact due to chronic loneliness (Hackel, Looser, & Van Bavel, 2014), when they have to cooperate with the entities (Hertz & Wiese, 2016) or when they have interacted with the entities for a long time. For instance, soldiers who interacted with search robots on an everyday basis in potentially dangerous situations have been reported to refuse to agree to install updates on their robots because they feared this would change their “personality” (Singer, 2008; Carpenter, 2016).

Mind perception based on physical features follows a categorical pattern with qualitative changes in mind ratings at the human-nonhuman *category boundary*, which is located at around 63% on a spectrum from 0% to 100% physical humanness (range: 55-70%; Looser & Wheatley, 2010; Cheetham, Suter, & Jancke, 2011, 2014; Hackel, Looser, & Van

Bavel, 2014; Martini, Gonzalez & Wiese, 2016). “Physical humanness” refers to the percentage of a human image contained in human-nonhuman morphed stimuli (e.g., a 30% human-robot morph contains 30% of a human and 70% of a robot image). Stimuli with significantly less (or more) than 63% physical humanness are easily categorized as non-human (or human) whereas stimuli straddling the categorical boundary are harder to categorize (i.e., *categorical identification*: human or nonhuman?) but easier to discriminate (i.e., *perceptual discrimination*: same or different?) than equally similar stimulus pairs located on the same side of the boundary (Repp, 1984; Harnad, 1987; Goldstone & Hendrickson, 2010; Cheetham et al., 2011, 2014). Thus, categorizing stimuli as human or nonhuman should be associated with increased cognitive load due to enhanced processing costs when their mind status is ambiguous because they fall on the human-nonhuman category boundary (i.e., longest reaction times and lowest hit rates; Cheetham et al., 2011, 2014). Cognitive performance on tasks incorporating such agents (e.g., interaction with humanoid avatars/robots) should be negatively impacted even if the agent’s mind status is not relevant to correct task performance (due to the bottom-up nature of mind perception).

Despite being an important question, experiments investigating the effect of mind perception on cognitive performance are surprisingly rare and have yielded mixed results (e.g., Teufel, Fletcher, & Davis, 2010; Wiese, Wykowska, Zwickel & Müller, 2012; Wykowska, Wiese, Prosser, & Müller, 2014; Wiese, Wykowska, & Müller, 2014; Martini, Buzzell, & Wiese, 2015; Abubshait & Wiese, 2017). For instance, when mind perception was manipulated *explicitly* via instruction (i.e., mind status of an agent was explicitly considered: low for a pre-programmed robot, high for a human-controlled robot), interactions with agents believed to have a mind resulted in better performance than interactions with agents be-

lieved to be mechanistic (e.g., Wiese et al., 2012; Wiese et al., 2014; Wykowska et al., 2014). However, when mind perception was manipulated *implicitly* via physical humanness (i.e., morphed agents of varying degrees of humanness were embedded in a cognitive task and their mind status was neither relevant to perform the task correctly, nor explicitly considered), decreased cognitive performance was observed during interactions with agents with ambiguous mind statuses (Martini et al., 2015), suggesting that efforts to implicitly categorize ambiguous agents led to cognitive conflict processing caused by interference of competing category representations, which drained cognitive resources from the main task. This interpretation is in line with established biased-competition models of visual processing (Desimone & Duncan, 1995), claiming that possible interpretations of ambiguous stimuli compete for representation in visual networks causing a cognitive conflict, and that cognitive resources are needed to direct selective attention to stimuli features that favor one explanation over the others (via top-down inhibition of the alternative category representations) in order to resolve the cognitive conflict (Meng, & Tong, 2004; Sterzer, Klein-schmidt, & Rees, 2009; Ferrey, Burleigh & Fenske, 2015). This suggests that (a) ambiguously human stimuli have the potential to induce a cognitive conflict due to co-activation of competing categories, which drains cognitive resources from a main task even if categorization is task-irrelevant, and (b) reduced amounts of cognitive resources available while performing a main task might not be sufficient to resolve potential cognitive conflict leading to sustained drainage of resources over time on task.

Aim of Study

While previous studies have shown THAT implicit mind perception has a negative effect on

cognitive performance, they did not investigate HOW mind perception causes this modulation in performance. Based in the studies outlined above, we hypothesize that implicit mind perception negatively impacts cognitive performance on tasks containing stimuli with ambiguous mind status due to category co-activation (human vs. robot), which leads to a cognitive conflict and drainage of cognitive resources over time. If that were true, embedding ambiguous stimuli into a task that is sensitive to the drainage of cognitive resources should more negatively impact task performance compared to stimuli that clearly belong in either the “human” or the “nonhuman” category. In line with the implicit nature of mind perception, the negative effect of an ambiguous mind status on cognitive performance should be observable, even if the agent’s category membership is irrelevant to the task (Palermo & Rhodes, 2007).

To investigate this question, we created a series of morphed stimuli of varying degrees of physical humanness (i.e., *human*: 100% human; *humanoid*: 70% human; *robotoid*: 30% human; *robot*: 0% human) and embedded them into a vigilance task known to be sensitive to the drainage of cognitive resources as indicated by a decline of performance over time on task (*vigilance decrement*; Mackworth, 1948; Davies & Parasuraman, 1982; Warm, Matthews, & Finomore, 2008; Warm, Parasuraman, & Matthews, 2008). In particular, participants performed a threat detection task that required them to pay sustained attention to the morphed stimuli as the stimuli directed their eye gaze either toward or away from threat-related probes (i.e., gun versus hairdryer; Parasuraman, de Visser, Clarke, McGarry, Hussey, Shaw, & Thompson, 2009). Participants were instructed to respond only when the stimuli looked toward the gun (i.e., critical trials), but not in other conditions (i.e., eyes directed toward hairdryer or away from either gun or hair dryer). Importantly, mind status

was not relevant to the task such that a response had to be given when *any kind of agent* looked at the gun and not specifically when a human-like agent did so. In order to qualify as a vigilance task, critical trials appeared with a low frequency and performance was measured in terms of *hit rates* (i.e., percentage of correct responses on critical trials), *false alarm rates* (i.e., percentage of responses on non-critical trials) and *reaction times* (i.e., reaction time to correct detections). If implicit mind perception caused a cognitive conflict for stimuli located at the categorical boundary between human and nonhuman (63% physical humanness), the vigilance decrement would be more pronounced for the humanoid stimulus (i.e., 70% physical humanness: ambiguous) than for the robot, robotoid and human stimuli (i.e., 0%, 30% and 100% physical humanness: non-ambiguous; **Experiment 1**). Furthermore, if a more pronounced vigilance decrement for the humanoid agent was due to cognitive conflict processing, facilitating categorization by pre-exposing participants to the stimuli before the actual vigilance task would attenuate vigilance decrements for the humanoid but not change vigilance performance for the other agents (**Experiments 2 & 3**). The assumption that pre-exposure would alleviate cognitive conflict for the humanoid agent was based on previous reports that conflict perception on one trial leads to transient upregulation of selective attention in the next trial, which then biases perception towards one particular interpretation and reduces cognitive conflict even without explicit instructions to do so (i.e., *conflict adaptation*; Gratton, Coles, & Donchin, 1992; Botvinick, Braver, Barch, Carter, & Cohen, 2001; Egner & Hirsch, 2006). By exposing participants to the stimuli outside of the vigilance task, we expected to speed up conflict adaptation by allowing them to resolve the conflict without being forced to concurrently perform another task that requires resources. If pre-exposure resolved cognitive conflict with the humanoid agent

prior to the vigilance task, competing activation of the “human” and “nonhuman” category should be reduced, which should free up resources for the vigilance task and reduce the decrement.

If ambiguity regarding the humanoid’s mind status had to be explicitly resolved prior to the vigilance task, rating the stimulus regarding different mind-related characteristics (e.g., capability to feel pain) during pre-exposure should attenuate the vigilance decrement (**Experiment 2**). However, if conflict adaptation took place implicitly (i.e., without explicit instructions to resolve categorization ambiguity), in-depth processing of the physical features of the stimuli (e.g., judging face shape) outside of the vigilance task should suffice to resolve categorical ambiguity and reduce the decrement for the humanoid agent (**Experiment 3**).

EXPERIMENT 1

Methods & Materials

Participants

Fifty-nine undergraduate students (47 female, 12 male; mean age: 20.3 years; SD: 4; range: 17-42 years) participated in Experiment 1. One outlier was removed (± 3 SD from the mean hit rate), leaving 58 participants who were randomly assigned to one of the following conditions: *human* (n=15, mean age: 21.1 years), *humanoid* (n=14, mean age: 19.1 years), *robotoid* (n=14, mean age: 20 years), and *robot* (n=15, mean age: 20.7 years). All participants gave informed consent, received course credit for participation and reported normal or corrected-to-normal visual acuity.

Apparatus

The stimuli were presented on a 19-inch monitor with the refresh rate set at 85 Hz. Reaction time and error rate measures were based on standard keyboard responses. Participants were seated approximately 57cm from the monitor and the experimenter ensured that participants were centered with respect to the monitor. The experiment was controlled by the software Experiment Builder (SR Research Ltd., Ontario, Canada); Fantamorph (Version 3; Abrosoft Co., Beijing China) was used to create the agent images.

Stimuli

Agents

The agents were designed to represent varying levels of physical humanness along a spectrum from robot to human. For that purpose, a human image (F07 from KDEF database; Lundqvist, Flykt, & Öhman, 1998) was morphed into a robot image (Meka robotics head S2) in increments of 10%. Four images out of this spectrum were chosen as stimuli for the current experiment: *robot* (0% human), *robotoid* (30% human), *humanoid* (70% human) and *human* (100% human); see **Figure 1**².

The agent stimuli were 16.24° x 16.24° in size, depicted in color on a white background and presented in full frontal orientation with eyes positioned on the central horizontal axis of the screen. For left- and rightward gaze, irises and pupils in the eyes were shifted with Photoshop to deviate 0.44° from direct gaze.

² Physical humanness refers to the percentage amount of the human image that is contained in the morphed image. For instance, a 60% human morph contains 40% of the Meka image and 60% of the F07 human image. Please note that although physical humanness is manipulated parametrically, perceived humanness follows a qualitative pattern (in line with Looser & Wheatley, 2010; Hackel et al., 2014; Martini et al., 2016).

Probes

Two different probes were used in this study: a gun and a hair dryer; see **Figure 2**. The probes were similar to those used in Parasuraman et al. (2009) and were selected for their morphologic similarities but different threat relevance (gun: threatening versus hair dryer: non-threatening). The probes measured 2.86° in width and 2.26° in height, were presented in black color on a white background and were located on the horizontal axis, 5.72° away from the center of the screen.

Procedure

In all conditions, the task was composed of 8 continuous blocks of 72 trials each (576 total trials) resulting in an overall time on task of 24 minutes and an event rate of 24 events per minute, in line with studies that have shown that vigilance decrements can be found in abbreviated versions of the classical paradigm (Temple, Warm, Dember, Jones, LaGrange, & Matthews, 2000). The blocks were purely used for analysis purposes and participants did not receive breaks during the task. Targets were presented at a rate of 9 times per block (9 out of 72 trials or 12.5% of all trials).

The sequence of events on a given trial is shown in **Figure 3**. At the beginning of each trial, a fixation cross was presented for 250 ms, followed by the presentation of the agent looking straight for 500 ms. Afterwards, the stimulus changed gaze direction, either looking to the left or the right side of the screen for 200 ms, followed by the probe stimulus appearing either at the gazed-at or gazed-away location for 100 ms. After that, the screen went blank allowing participants to respond: On critical trials, participants had to press the

spacebar to indicate probe detection (i.e., agent looked at the gun), while on non-critical trials (i.e., agent looked at hair dryer or looked away from the probe), participants did not press any button. If no answer was given, the next trial started after 1200 ms. Responses occurring after the 1200ms response window were classified as false alarms.

At the beginning of the experiment, participants gave informed consent and were introduced to the task via an instructional presentation that they could click through at their own pace. Afterwards, they had to pass a 3-minute practice session that mirrored the experimental procedure with the addition of auditory feedback (i.e., chime sound for correct responses; buzzing sound for incorrect or missed responses). Participants had to perform at least at 66% accuracy in the practice to be able to advance to the actual task, which is a commonly used threshold for practice trials in vigilance paradigms (e.g., Molloy & Parasuraman, 1996; Shaw et al., 2013). Participants who did not meet the threshold after the first attempt were permitted to take the practice again and were excluded from the study if they did not manage to meet the threshold after the second attempt. The threshold was enforced to ensure that participants understood the task and that any detrimental effects on performance during the experiment would not be due to misunderstanding the task. The experimental task did not have auditory feedback due to its potential alerting effect (Manly, Hawkins, Evans, Woldt, & Robertson, 2002).

Results

33 individuals were unable to meet the practice requirements and were disqualified from the study. Of the 58 qualified participants (79% female), 15 completed the practice session twice (73% female). The overall vigilance performance of those who repeated practice did

not differ from those who took practice only once according to hit rate (once: $M = 94.2\%$, $SD = 8.3\%$; twice: $M = 91.8\%$, $SD = 10.1\%$), reaction time (once: $M_{rt} = 543.6$ ms, $SD = 80$; twice: 521.9 ms, $SD = 103.9$), or false alarms (once: 1.7% , $SD = 1.8\%$; twice: 3.1% , $SD = 2.7\%$).

The performance data from the vigilance task was analyzed using a 4 x 8 mixed design ANOVA, with the between-participants factor *Agent Type* (human, humanoid, robotoid, robot) and the within-participants factor *Block* (1 to 8). As dependent variables, hit rate (i.e., percentage of correct responses to critical signals), false alarm rate (i.e., percentage of responses given when critical signal was not present), and response time (i.e., reaction time to correct detections) were measured. Nonparametric signal detection measures of sensitivity (A') and response criterion ($B''d$) were calculated and analyzed using a comparable 4 x 8 mixed ANOVA. The purpose of this analysis was to examine the cause of any observed performance decrement (i.e., do changes in performance reflect a shift in the response criterion, perceptual sensitivity, or both?). Additionally, in order to examine how performance changed over time, we calculated slopes for each agent by averaging the individual performance (i.e., hit rate, false alarm, reaction time) of each participant with a given agent across blocks ($y = mx + b$, with " m " representing the slope and " b " representing the intercept). The performance slopes were then compared across agents using a between-participants ANOVA with agent type as the independent variable. Greenhouse-Geisser corrections were used to correct for analyses that violated the assumption of sphericity. All post-hoc tests were Bonferroni-corrected for multiple comparisons.

Hit Rates

Descriptive statistics of the analysis of the hit rates across agents and blocks are shown in

Table 1. The 4 x 8 mixed-ANOVA on hit rates revealed a main effect of *Block* ($F(3.4, 185.6) = 4.88, p = .001, \eta^2 = .08$), indicating that there was a significant decrease in hit rates over time. In addition, probe detection on critical trials differed between agents, as indicated by a significant main effect of *Agent Type* ($F(3, 54) = 3.75, p = .016, \eta^2 = .17$). Hit rates were significantly lower for the humanoid agent ($M = 87.3\%, SD = 12.17$) compared to the robot, robotoid, and human conditions combined ($M = 95.52\%, SD = 6.42; F(1, 54) = 10.67, p = .002, \eta^2 = .17$).

Most importantly, there was a significant interaction effect between *Block* and *Agent Type* ($F(10.3, 185.6) = 2.576, p = .006, \eta^2 = .125$), indicating that the degree of the vigilance decrement varied significantly across agents. In particular, hit rates in the humanoid condition decreased dramatically over time (-27.8%, i.e. from 100% in block 1 to 72.2% in block 8), while vigilance decrements in the other conditions were comparatively small (range: -3% for the robot condition, -3.3% for the robotoid condition, and +4.5% for the human condition). Tests of simple effects indicated that while the humanoid condition had a significantly higher hit rate in the first block compared to the human condition ($p < .05$), this effect reversed by the end of the task and the humanoid condition's hit rate was significantly lower than all other conditions ($p < .01$). In addition, the ANOVA on average hit rate slope showed a significant main effect of *Agent Type* ($F(3, 54) = 5.67, p = .002, \eta^2 = .24$), and planned post-hoc contrasts revealed a steeper and more negative decrement in the humanoid condition ($M = -3.63, SD = 4.4$) compared to the other conditions combined ($M = -.37, SD = 1.88; F(1, 54) = 15.4, p < .001, \eta^2 = .22$); see **Figure 4a**.

False Alarms

Descriptive statistics of the analysis of the false alarm rates across agents and blocks are shown in **Table 2**. The 4 x 8 mixed-ANOVA on false alarm rates revealed a significant main effect of *Block* ($F(2.7, 147.7) = 6.47, p = .001, \eta^2 = .11$), with a slight decrease in false alarm rates over time across all conditions. Furthermore, there was a main effect of *Agent Type* ($F(3,54) = 3.82, p = .015, \eta^2 = .18$), with larger false alarm rates for the humanoid condition ($M = 3.02\%, SD = 2.66$) compared to the robot, robotoid, and human conditions combined ($M = 1.7\%, SD = 1.82; F(1,54) = 4.90, p = .031, \eta^2 = .08$). There was no significant difference in vigilance performance across agents over time, as indicated by a non-significant interaction effect between *Agent Type* and *Block* ($F(8.2, 147.7) = 1.0, p = .439, \eta^2 = .05$) and a non-significant main effect of *Agent Type* ($F(3,54) = .376, p = .771, \eta^2 = .02$) on false alarm slopes; see **Figure 4b**.

Reaction Times

Descriptive statistics of the analysis of reaction times across agents and blocks are shown in **Table 3**. The 4 x 8 mixed-ANOVA on reaction times for correct detections revealed a significant main effect of *Block* ($F(4.7, 251.2) = 13.58, p < .01, \eta^2 = .20$), indicating a significant increase in reaction times over time across conditions. The main effect of *Agent Type* ($F(3, 54) = 2, p = .125, \eta^2 = .1$), the interaction between *Agent Type* and *Block* ($F(14.3, 251.2) = 1.01, p = .44, \eta^2 = .05$) and the ANOVA examining the effect of *Agent Type* on reaction time slopes ($F(3,54) = .23, p = .88, \eta^2 = .01$) did not reveal significant effects; see **Figure 4c**.

Signal Detection Analysis

An analysis of signal detection measures reveals a main effect of block on the response criterion ($B''d$; $F(7,378) = 14.915, p < .001, \eta^2 = .216$). Pairwise comparisons indicated that the response criterion in blocks 1 and 2 was significantly lower than in blocks 4-8, suggesting that the response criterion became more conservative over time on task. The main effect of *Agent Type* ($F(3,54) = 1.565, p = .208, \eta^2 = .08$) and the interaction between *Block* and *Agent Type* ($F(1.45,314.22) = 1.324, p = .173, \eta^2 = .069$) were not significant, indicating that the average response criterion and its change over time were not affected by the agent type.

While sensitivity (A') did not change over time on task (i.e., insignificant main effect of *Block*, $F(2.68,144.7) = 2.158, p = .103, \eta^2 = .038$), there was a main effect of *Agent Type* on sensitivity ($F(3,54) = 4.3, p = .009, \eta^2 = .193$) such that sensitivity was significantly lower in the humanoid condition ($M = .95, SD = .04$) than the robot ($M = .97, SD = .02; p < .05$) and robotoid conditions ($M = .98, SD = .01; p < .05$). In addition, there was a significant *Block x Agent Type* interaction ($F(8.04,144.7) = 2.569, p = .012, \eta^2 = .125$), indicating that changes in sensitivity over time differed by agent type. Pairwise comparisons revealed that the interaction was driven by a significantly stronger decrease in sensitivity for the humanoid than all other agents on block 8 ($p < .05$).

In summary, conservatism increased over time with all agent types. Sensitivity was stable throughout the task for all agent types except the humanoid, which exhibited a significant decrease in sensitivity at the end of the task that parallels the humanoid-specific hit rate decrement.

Discussion

The results of Experiment 1 showed the expected vigilance decrement, as reflected in the negative percent slope change for hits and the positive slope change for reaction times across all agents. In addition, the response criterion increased over time on task, a result consistent with typical findings in vigilance studies (Parasuraman, 1979). No vigilance decrement was found in the analysis of the false alarm slopes, which is also seen in other vigilance experiments (Parasuraman & Mouloua, 1987). It should be pointed out, however, that the average false alarms rate was low across all conditions, and the sensitivity index ($d' = 3.56$) indicates that overall, the task can be considered moderately easy (cf. Craig, 1984).

There was a general effect of agent type on vigilance performance, with lower hit rates, higher false alarm rates, and lower sensitivity for the humanoid agent compared to the other agents (specifically, lower than the robot and robotoid for the measure of sensitivity). Importantly, agent type also affected vigilance performance over time, with a significantly stronger vigilance decrement in both performance and sensitivity for the humanoid agent compared to the other three agents. According to cognitive resource theory (Wickens & Kessel, 1980; Parasuraman & Mouloua, 1987), performance decrements result from the depletion of cognitive resources over time on task with larger decrements associated with the use of more cognitive resources. Hence, the observation of greater performance and sensitivity decrements in the humanoid condition suggests that processing the humanoid agent was cognitively more demanding than the processing of any other agent. This is in line with the hypothesis that stimuli located at the human-nonhuman category boundary drain cognitive resources over time (in line with the assumption of cognitive conflict processing) and negatively impact performance on sustained attention tasks. The presence of a

category boundary at around 70% physical humanness is also in line with previous studies (Looser & Wheatley, 2010; Cheetham et al., 2011, 2014; Hackel et al., 2014; Martini, Gonzalez & Wiese, 2016).

In order to test more directly if the stimuli used in the current experiment induced a cognitive conflict with regard to mind perception, we asked 55 participants (24 male; mean age= 36.8 years) in a post-hoc experiment to rate the stimuli regarding their ability of having a mind (i.e., using a Likert scale, following the procedure of Martini et al., 2016) and 161 participants (78 male; mean age= 33.8 years) in another post-hoc experiment to categorize the stimuli in a force-choice scenario into human and nonhuman (i.e., measuring percent of human versus non-human categorizations and reaction times, following the procedure of Lakatos, Gácsi, Konok, Brúder, Bereczky, Korondi, & Miklósi, 2014). Both experiments revealed that the humanoid stimulus was more difficult to classify than the other stimuli, as indicated by more ambiguous mind ratings, as well as less consensus and longer reaction times for the categorization of the humanoid stimulus compared to the other three stimuli; see *Supplementary Materials (Section 1)* for more details on the methods used and **Figure 5** for a summary of the results.

The results of Experiment 1 support the hypothesis that vigilance performance is significantly impaired for the humanoid stimulus due to the drainage of cognitive resources over time, which seems to be the result of cognitive conflict processing induced by stimuli with ambiguous mind status. The results also indicate that mind perception seems to be hard to suppress even if the category membership of the stimulus is irrelevant for the task and processing its mind status is detrimental to cognitive performance. If increased vigilance decrements for the humanoid agent were indeed due to a cognitive conflict, then giv-

ing participants the chance to resolve the conflict before they perform the vigilance task should significantly recover performance for the humanoid stimulus, but not impact performance for the robot, robotoid and human stimulus. Experiment 2 investigates this hypothesis by pre-exposing participants to all stimuli before the actual vigilance task.

EXPERIMENT 2

Experiment 2 explored whether pre-exposing participants to the agent stimuli before they perform the vigilance task would diminish the cognitive conflict caused by the humanoid agent in Experiment 1. We hypothesized that if the performance decrement in Experiment 1 was caused by a cognitive conflict, then pre-exposing participants to the agents before the vigilance task should resolve any cognitive ambiguity (due to increased time and depth of processing) and thus lead to a reduction in the vigilance decrement for the humanoid stimulus. In order to assure that participants explicitly resolve the cognitive conflict, the stimuli were embedded in a social assessment task that required judgments regarding the stimuli's ability of having a mind, and should resolve any uncertainty regarding whether the stimulus should be treated as an agent 'with a mind' or 'without a mind'.

Methods & Materials

Participants

Forty-four undergraduate students (23 female, 21 male; mean age: 20.5 years; SD: 2.3) participated in Experiment 2 for course credit. One outlier was removed (± 3 SD from the mean hit rate); the remaining 43 were randomly assigned to one of the following conditions: *human* (n=14, mean age: 20.6 years), *humanoid* (n=15, mean age: 21.5 years), and

robot (n=15, mean age: 19.3 years). The robotoid condition of Experiment 1 was not included in Experiment 2 given that it did not show increased vigilance decrements in Experiment 1. All participants gave informed consent, received course credit for their participation and reported normal or corrected-to-normal visual acuity.

Apparatus & Stimuli

Experiment 2 was comparable to Experiment 1, with two exceptions: First, instead of four agents, we presented only three agents, namely the *human* (100% human), the *humanoid* (70% human) and the *robot* (0% human). This change was made to increase the efficiency of the paradigm since there were no significant differences between the robot and the robotoid condition in Experiment 1 on any measure (i.e., both clearly fall into the nonhuman category). This change should not have any impact on participants' performance on the vigilance task given that agent type is manipulated between participants in both experiments. Second, participants were pre-exposed to the agent stimuli during a social judgment task where they were asked to rate the stimuli regarding their capability of having a mind using the items from Martini et al. (2016), which are similar to those used previously (Weisman et al., 2017; Hackel et al., 2014; Looser & Wheatley, 2010). In particular, the questions pertained to being alive, having a mind, feeling pain, feeling the need to hang out, and being interested in social interactions (see **Table 4**). This was done because we wanted to assure that (1) participants actively processed the images in sufficient depth, and (2) explicitly resolved the cognitive conflict during pre-exposure by making judgments about the agents' mind status. The social judgment task took about 10 minutes and was administered right

before the vigilance task. Pre-exposure was done for all agent conditions so that no selective effects of this manipulation on the humanoid condition were to be expected.

Procedure

The experimental procedure was comparable to Experiment 1, with the exception of the social judgment task being administered before the vigilance task.

At the beginning of each session, participants gave informed consent and were instructed to complete a survey during which they would rate a series of agent images (i.e., social judgment task). During the task, the same agent images that were used for the vigilance task were presented (one image at a time, order randomized, faces looking straight) and participants were asked to rate the images regarding their capability of 'having a mind' on five different questions (**Table 4**). The questions were presented at the same time as the agent images, were depicted underneath the images and appeared together with the agent images. Participants rated the agents on all five questions using a 7-point Likert scale (1= Definitely Not; 7= Definitely). Images, items and rating scales remained on the screen until a response was given (i.e., there was no time limit on the task). Upon completing the ratings for one agent, the task advanced to the next agent; the order in which the agents were presented during the rating task was randomized and the rating procedure was repeated until all agents were rated. Afterwards, the experiment continued with the same vigilance task as in Experiment 1.

Similar to Experiment 1, the vigilance task was performed with only one agent and agent type was randomly assigned to participants. The social judgment task and vigilance task together took a total of 45 minutes to complete.

Analysis

The analysis of the vigilance data in Experiment 2 focuses on hit rates (i.e., percentage of correct responses to critical signals) given that hit rate was the variable where the significant interaction between block and agent type was found in Experiment 1; the analyses of false alarm rates, and response times are comparable to Experiment 1 and are reported in the *Supplementary Materials*. Similar to Experiment 1, the hit rate data was first analyzed using a mixed 3 x 8 ANOVA with *Agent Type* (robot, humanoid, human) as between-factor and *Block* (1 to 8) as within-factor. This was followed by a univariate ANOVA on hit rate slopes with the between-factor *Agent Type* (robot, humanoid, human) to examine performance changes over time. For all analyses, the Greenhouse-Geisser correction is used to correct for analyses that violated the assumption of sphericity and all post-hoc tests are Bonferroni-corrected for multiple comparisons.

Results

23 students did not pass practice. Of the remaining 43 participants (50% female), 8 took practice twice (78% female). Vigilance performance for those who took practice once and those who took practice twice did not differ by hit rate (once: $M = 95.8\%$, $SD = 4.6\%$; twice: $M = 95.7\%$, $SD = 5.7\%$), reaction time (once: $M = 568.7$ ms, $SD = 93.4$; twice: $M = 538.6$ ms, $SD = 67$), or false alarm rate (once: $M = 1.7\%$, $SD = 2.5\%$; twice: $M = 2.1\%$, $SD = 3.5\%$).

Social Judgment Task

The social judgment task results followed a dual-linear pattern, with the robot receiving low mind ratings ($M = 2.05$, $SD = 1.18$), and the human agent receiving high mind ratings

($M = 6.6$, $SD = .83$); the humanoid agent received intermediate mind ratings ($M = 4.02$, $SD = 1.57$), indicating that the perception of its ability of having a mind was ambiguous. The data was not analyzed further since the main purpose of this task was to make sure that participants process the stimuli in sufficient depth and explicitly resolve potential cognitive conflicts associated with the agents' mind status.

Hit Rates

Descriptive statistics of the analysis of the hit rates across agents and blocks are shown in **Table 5**. The 3 x 8 mixed-ANOVA on hit rates revealed a significant main effect of *Block* ($F(4.6, 190) = 2.83$, $p = .02$, $\eta^2 = .07$), indicating that there was a significant decrease in hit rates over time; the main effect of *Agent Type* was not significant ($F(2, 41) = 1.55$, $p = .22$, $\eta^2 = .07$). Most importantly, in contrast to Experiment 1 the interaction effect of *Agent Type* and *Block* on hit rates ($F(9.3, 190) = 1.09$, $p = .37$, $\eta^2 = .05$), and main effect of *Agent Type* on hit rate slopes ($F(2, 41) = 1.22$, $p = .31$, $\eta^2 = .06$) were not significant, indicating that the vigilance decrement that was observed for the humanoid agent in Experiment 1 (**Figure 6a**) vanished when participants were pre-exposed to the agents before performing the vigilance task, with no significant difference in vigilance performance between agent types after pre-exposure (**Figure 6b**).

Signal Detection Analysis

As in Experiment 1, the response criterion increased over time on task according to a significant main effect of *Block* ($F(4.74, 194.25) = 8.687$, $p < .001$, $\eta^2 = .175$) and pairwise comparisons demonstrating that the response criterion in block 1 is significantly lower than in

blocks 4-8 ($p < .05$ for all comparisons). This reflects an increasingly conservative criterion. The main effect of *Block* on sensitivity ($F(3.96, 162.33) = 1.397, p = .238, \eta^2 = .033$), main effects of *Agent Type* on sensitivity ($F(2, 41) = 1.425, p = .252, \eta^2 = .065$) and response criterion ($F(2, 41) = .782, p = .464, \eta^2 = .037$), and interaction between *Block* and *Agent Type* on sensitivity ($F(7.92, 162.33) = 1.014, p = .427, \eta^2 = .047$) and response criterion ($F(9.48, 194.25) = .540, p = .852, \eta^2 = .026$) were all non-significant.

Discussion

The purpose of Experiment 2 was to test whether the vigilance decrement for the humanoid agent in Experiment 1 was due to cognitive conflict processing and would therefore vanish if participants got the chance to familiarize themselves with the stimuli (i.e., process them in sufficient depth) before performing the actual vigilance task. In order to investigate this hypothesis, participants performed a social judgment task prior to the vigilance task during which they rated the different stimuli regarding their capabilities of having a mind (i.e., explicit mind perception). The results show a significant vigilance decrement for all agents as indicated by a decrease in hit rates over time (i.e., main effect of block on hit rates) and increasingly conservative response criterion; however, while the vigilance decrements for the human and the robot agent were unchanged compared to Experiment 1, the vigilance decrement for the humanoid agent was reduced and did not significantly differ anymore from the decrements induced by the human and robot agents.

The results support the hypothesis that the observed vigilance decrement for the humanoid agent in Experiment 1 was based on implicit categorization processes ('having a mind' versus 'not having a mind') that caused a cognitive conflict associated with increased

workload due to the stimulus' proximity to the human-nonhuman category boundary. Allowing participants to process and judge the stimuli before the vigilance task enabled them to explicitly reduce the cognitive conflict and lead to increased performance and sensitivity on the vigilance task with the humanoid agent (i.e., reduced vigilance decrement). As expected, pre-exposure reduced the vigilance decrement for the humanoid agent (i.e. decrement was no longer steeper than the other agents), indicating that the social judgment task did not improve performance on the vigilance task in general, but exclusively for the ambiguous stimulus (i.e., located at the category boundary between human and non-human).

The question remains, however, whether the positive effect of pre-exposure on performance was simply due to increased processing time and depth (i.e., prolonged perceptual processing time is sufficient to reduce cognitive conflict: *implicit mind perception*) or due to the nature of the social judgment task (i.e., making judgments about internal states: *explicit mind perception*). More specifically, being asked to judge the agents' internal states forced participants to make explicit categorical decisions as to whether they want to treat the stimulus as human or non-human AND allowed them to process the stimuli more deeply, which makes it impossible to judge which of the two features of the social judgment task caused the vigilance decrement to lessen. Experiment 3 addresses this issue by examining whether questioning that is not related to mind perception but allows participants to process the stimuli in depth (i.e., judging the stimuli in terms of physical features, such as face shape, instead of internal states) still has a positive effect on vigilance performance for the humanoid agent (i.e., significant decrease compared to Experiment 1).

EXPERIMENT 3

Experiment 3 examines whether the positive effect of pre-exposure to the humanoid agent on vigilance performance requires explicit judgments of its mind status or whether it can also be observed when participants judge the stimulus regarding physical features that are unrelated to its mind status but allow that the stimulus is processed in depth. To investigate this question, we repeated the procedure of Experiment 2, but this time asked participants to rate the stimuli regarding physical features that were not directly related to mind perception (e.g., face shape). If mere exposure (i.e., without explicit judgments regarding the mind status) is sufficient for a reduction of cognitive conflict processing associated with the humanoid agent, then a reduction in the vigilance decrement similar to Experiment 2 should also be noticeable in Experiment 3 (compared to Experiment 1). However, if explicit judgments about internal states are necessary for a reduction of the cognitive conflict, then the familiarization phase of Experiment 3 should not be sufficient to reduce vigilance decrements significantly.

Methods & Materials

Participants

44 undergraduate students were recruited to participate in Experiment 3 (28 female, 16 male; mean age: 19.8 years; SD: 2) and were randomly assigned to one of the following conditions: *human* (n=15, mean age: 19.1 years), *humanoid* (n=15, mean age: 20 years), and *robot* (n=14, mean age: 20.2 years). All participants gave informed consent, received course credit in exchange for their participation and reported normal or corrected-to-normal visual acuity.

Apparatus, Stimuli, Procedure & Analysis

The apparatus, stimuli, procedure and analysis for Experiment 3 were the same as in Experiment 2, with one exception: Instead of engaging in a social judgment task, participants performed a perceptual judgment task, during which they rated certain physical features of the stimuli. In particular, the questions pertained to age, attractiveness, face shape, gender, and hair length (see **Table 6** for items developed for this experiment). Similar to the social judgment task, the perceptual judgment task took about 10 minutes, was administered before the vigilance task and answers were given on a 7-point rating scale. Participants were allowed to attempt practice 3 times before disqualification.

Results

14 individuals were not able to pass practice. Of the 44 participants who did pass practice, 7 completed practice 2 times and 2 completed practice 3 times. The overall performance of these three groups was similar for hit rates (once: $M = 96.7\%$, $SD = 5.1\%$; twice: $M = 97.8\%$, $SD = 4\%$; three times: $M = 95.8\%$, $SD = 2\%$), false alarms (once: $M = 1\%$, $SD = 1.2\%$; twice: $M = .5\%$, $SD = .7\%$; three times: $M = 1.2\%$, $SD = 0\%$), and reaction times (once: $M = 555.6$ ms, $SD = 89$; twice: $M = 525.6$ ms, $SD = 81.2$; three times: $M = 558$ ms, $SD = 36.7$).

Perceptual Judgment Task

The perceptual judgment task required participants to rate physical, rather than mind status related attributes of the agents. Please note that although features like attractiveness or gender are more “social” in nature than shape or hair length, judging them does not explicitly require mind perception (i.e., stimuli without mind can be perceived as being of a cer-

tain gender or attractiveness: robots can look male or female and words can sound male or female; cats or dogs can be perceived as attractive and cockroaches or naked mole rats as unattractive), and to the best of our knowledge, questions regarding gender or attractiveness have not been used to examine mind perception (Gray et al., 2007; Waytz et al., 2010). In contrast to Experiment 2, where the mind rating was used as a manipulation-check, the perceptual rating in Experiment 3 was used only to assure that participants processed the faces during pre-exposure. Consequently, the data was not included in the main analysis, though summative data can be found in *Supplementary Materials*.

Hit Rates

Descriptive statistics of the analysis of the hit rates across agent type and block are shown in **Table 7**. The 3 x 8 mixed ANOVA on hit rates revealed non-significant main effects of *Block* ($F(5, 204) = 1.32, p = .26, \eta^2 = .03$) and *Agent Type* ($F(2, 41) = 1.52, p = .23, \eta^2 = .07$), a non-significant interaction between *Block* and *Agent Type* on hit rates ($F(10, 204) = 1.15, p = .33, \eta^2 = .05$), and a non-significant main effect of *Agent Type* on hit rate slopes ($F(2, 41) = 2.02, p = .15, \eta^2 = .09$). These results indicate that there was no vigilance decrement for any of the agents after the perceptual rating task (**Figure 6c**).

Signal Detection Analysis

Again, a significant main effect of *Block* on response criterion ($F(4.72, 193.69) = 3.166, p = .01, \eta^2 = .072$) and pairwise comparison demonstrating a significantly greater response criterion in block 8 than block 1 ($p < .05$) indicated that the response criterion increased over time on task (i.e., criterion becomes more conservative). The main effect of *Block* on sensitivity ($F(4.28, 175.29) = .846, p = .55, \eta^2 = .02$), as well as main effects of *Agent Type* on sen-

sitivity ($F(2,41) = 1.744, p = .188, \eta^2 = .078$) and response bias ($F(2,41) = .202, p = .818, \eta^2 = .01$), and interaction between *Block* and *Agent Type* on sensitivity ($F(8.55,175.29) = 1.408, p = .192, \eta^2 = .064$) and response criterion ($F(9.44,193.69) = .982, p = .459, \eta^2 = .046$) were non-significant.

Comparison across Experiments

In order to determine whether pre-exposing participants to the stimuli before the vigilance task significantly reduced the vigilance decrement, a 3 (Experiment) x 3 (Agent Type) ANOVA was conducted with hit rate slopes as the dependent variable. This analysis revealed significant main effects of *Experiment* ($F(2, 123) = 3.82, p = .03, \eta^2 = .06$) and *Agent Type* ($F(2, 123) = 8.88, p < .001, \eta^2 = .13$), as well as a significant interaction between *Experiment* and *Agent Type* ($F(4, 123) = 3.15, p = .02, \eta^2 = .09$). A post-hoc planned comparison revealed a significantly greater decrement for the humanoid agent in Experiment 1 than in Experiments 2 and 3 ($F(1, 41) = 9.98, p = .003, \eta^2 = .2$), with no significant difference between Experiments 2 and 3 ($F(1, 41) = .26, p = .62, \eta^2 = .01$). This shows that pre-exposure to the humanoid agent prior to the vigilance task significantly reduced the vigilance decrement induced by the humanoid agent, but that the kind of exposure (i.e., explicit mind perception in Experiment 2 versus implicit mind perception in Experiment 3) did not have an impact on the vigilance decrement (**Figure 6**).

Discussion

The purpose of Experiment 3 was to examine whether the reduction of the vigilance decrement for the humanoid agent observed in Experiment 2 was (a) simply due to increased

depth and length of processing in general, or (b) specifically due to the nature of the pre-exposure task that required making explicit judgments about the agents' internal states. In order to address this question, Experiment 3 pre-exposed participants to the agent stimuli before performing the vigilance task, but instead of asking for explicit mind judgments during pre-exposure, participants had to assess certain physical features of the stimuli, which allowed for more pronounced processing, but required no direct assessment of the agents' mind status (i.e., implicit mind perception).

The results of Experiment 3 followed the same pattern as Experiment 2 except with the elimination of the vigilance decrement for all agents including the humanoid after pre-exposure and, more notably, with a significant reduction of the vigilance decrement for the humanoid agent compared to Experiment 1. The positive effect of pre-exposure is therefore not specific to the assessment of mind-related agent features but results from a deeper processing of the perceptual features of the stimulus. Altogether, these findings suggest that pre-exposure does not need to invite explicit mind perception processes in order to resolve the cognitive conflict associated with the humanoid agent, but needs to simply allow for a deeper processing of the stimulus.

GENERAL DISCUSSION

The goal of the current study was to investigate whether mind perception has a negative impact on performance during tasks involving stimuli with ambiguous mind status (i.e., located at the human-nonhuman category boundary). The negative effect of ambiguous stimuli on performance was assumed to be due to a cognitive conflict as to whether to categorize them as human versus nonhuman, associated with increased workload and drainage of

cognitive resources over time in order to resolve the conflict. To investigate this question, we embedded morphed images of varying degrees of physical humanness into a vigilance task, which is known to be sensitive to the drainage of cognitive resources (e.g. Mackworth, 1948; Davies & Parasuraman, 1982; Warm et al., 2008). Following the cognitive conflict hypothesis, we expected to find stronger vigilance decrements for humanoid stimuli (i.e., 70% physical humanness) than for stimuli that can easily be categorized as either human or nonhuman. We expected the negative effect of categorization ambiguity on performance to vanish when participants were pre-exposed to the stimuli before the vigilance task either due to explicit processes of mind perception (i.e., agents were rated according to their ability of 'having a mind'; **Experiment 2**) or implicit mind perception processes associated with increased processing time and depth (i.e., agents were rated according to perceptual details of their appearance; **Experiment 3**).

The results show significant vigilance decrements across all agent types for hit rate (in Experiments 1 and 2) and response criterion (Experiments 1, 2, and 3), with no general effect of physical humanness on vigilance performance over time. However, without pre-exposure, the humanoid agent caused a significantly stronger vigilance decrement than the robot, robotoid and human agent in Experiment 1. According to cognitive resource theory (Wickens & Kessel, 1980; Parasuraman & Mouloua, 1987), this finding indicates that the humanoid agent was associated with a significantly stronger depletion of cognitive resources over time than any of the other agents, which is in line with our hypothesis that stimuli located at the categorical boundary cause a cognitive conflict that is associated with increased cognitive processing costs. The results are in line with previous studies that have shown a categorical boundary along the human-nonhuman spectrum (Looser & Wheatley,

2010; Cheetham et al., 2011, 2014; Hackel et al., 2014; Martini et al., 2016), associated with longer reaction times and higher ambiguity for stimuli located at the category boundary (e.g., Cheetham et al., 2011, 2014). They are also consistent with a Mouse-Tracking (which has been shown to be a reliable measure for cognitive conflict processing; Freeman & Ambady, 2010) study by Weis and Wiese (2017) that used the same stimuli as in the present experiment and showed that cognitive conflict processing was maximal (as measured by various parameters associated with the curvature of the Mouse movement) at 72% physical humanness.

The finding that proximity to the categorical boundary had a negative effect on vigilance performance even though categorization of the agents as human versus non-human was irrelevant to the task is in line with previous studies showing that mind perception is a highly automatic, bottom-up process (Palermo & Rhodes, 2007; Looser & Wheatley, 2010), and negatively impacts performance for ambiguous stimuli (Martini et al., 2015). Specifically, Martini and colleagues (2015) used the same morphs as in the current study in a counterpredictive gaze-cueing paradigm, where target letters appeared more likely at locations that the stimuli looked away from (i.e., uncued; 80% of the trials) than looked towards (i.e., cued; 20% of the trials), and participants had to respond to the identity of the target letter (i.e., F or T) as fast and accurately as possible. To perform the task well, participants needed to suppress reflexive shifts of attention to cued (unlikely) location and perform voluntary shifts of attention to the uncued (likely) location, which requires top-down control of attentional orienting associated with increased cognitive load (Friesen, Ristic, & Kingstone, 2004). In line with the current experiment, the results show that participants were significantly less able to suppress reflexive shifts of attention to the gaze-cued/unlikely target lo-

cation and to shift their attention to the uncued/likely target location when the gazer was the humanoid agent compared to when it was a 0% human, 40% human or 100% human agent, indicating that fewer cognitive resources were available in the humanoid condition to top-down control attentional orienting.

Pre-exposing participants to the stimuli prior to the vigilance task in Experiment 2 and 3 significantly attenuated the sensitivity decrement for the humanoid agent while not changing the decrements for the other agents. This selective attenuation of the humanoid decrement was also observed for accuracy in Experiment 2 (but not Experiment 3). Crucially, whether participants rated the stimuli regarding their capacity of having internal states (Experiment 2: explicit mind perception) or regarding general physical features (Experiment 3: implicit mind perception), both types of pre-exposure significantly attenuated the sensitivity decrement for the humanoid agent compared to Experiment 1. This shows that explicitly instructing participants to process the perceptual features of the stimuli in detail was sufficient to alleviate the cognitive conflict it induces (see Bornstein & D'Agostino, 1992; Lee, 2001), and that explicit processes of mind perception were not required to reduce the vigilance decrement. Hence, it seems that cognitive conflict associated with mind perception is triggered by bottom-up mechanisms related to the processing of ambiguous perceptual features (Looser & Wheatley, 2010; Gao et al., 2010) and automatic co-activation of competing category representations (Ferrey et al., 2015) can only be resolved if alternative categorical explanations are inhibited by focusing selective attention on those perceptual features that support the chosen categorization (Desimone & Duncan, 1995). This shows that being able to process the perceptual details of the stimuli is necessary to top-down control the inhibition of alternative representations, while the mere familiariza-

tion with the stimuli during the vigilance task is not sufficient to reduce cognitive conflict processing. Please note, however, that we cannot ultimately determine whether mere exposure to the stimuli outside of the vigilance task (i.e., without instructions that required processing of perceptual details) would have yielded a similar reduction in vigilance decrement, since the current study did not contain a passive viewing condition (i.e., stimuli would be viewed during pre-exposure but not rated or classified).

There are several other theories that could potentially explain the current findings, besides the cognitive conflict hypothesis. First, participants might find it unpleasant to look at stimuli that have a high degree of human-likeness but are not perfectly human, which is in line with the *uncanny valley hypothesis* (Mori, MacDorman, & Kageki, 2012): stimuli that are very human-like, but not perfectly human are assumed to elicit more negative affinity (i.e., lower familiarity and more negative emotional valence) than other stimuli that are either clearly machines or humans. The negative affinity associated with ambiguous stimuli is thought to increase the observers' energetic arousal, leading to decreased performance by narrowing the attentional focus (Kätsyri, Förger, Mäkäraäinen, & Takala, 2015). To address this, we would argue that although being frequently cited in the literature, recent reviews come to the conclusion that empirical evidence for the uncanny valley (and fact that it is associated with increases energetic arousal) is either lacking or inconsistent (Pollick, 2009; Zlotowski, Proudfoot, & Bartneck, 2013; Kätsyri et al., 2015). In addition, while increased energetic arousal has been shown to have negative effects on cognitive performance in general (e.g., Eysenck & Calvo, 1992), performance on vigilance tasks seems to benefit from increased energetic arousal (Parasuraman, Warm, & See, 1998), making the

uncanny valley hypothesis an unlikely explanation for our findings (i.e., one would expect the humanoid stimulus to have the smallest vigilance decrement in Experiment 1).

Second, previous studies have shown that ambiguous stimuli might be perceived as threat that captures participants' attention, and makes it harder for them to disengage and focus their attention elsewhere (i.e., *delayed disengagement*; Fox, Russo, & Dutton, 2002), which can have negative effects on cognitive performance because attentional orienting to task-relevant information in the environment is delayed (Fox, Russo, Bowles, & Dutton, 2001; Belopolsky, Devue, & Theeuwes, 2011). To the best of our knowledge, there is no empirical evidence that ambiguous faces induce delayed disengagement due to automatic threat detection. In fact, Smith and Wiese (2016) have shown that while processing face-like stimuli caused delayed disengagement effects in general, the effect was not stronger for humanoid faces than for human and robot stimuli, making the threat detection hypothesis an unlikely explanation for our findings (i.e., one would expect no difference in vigilance decrements between agents in Experiment 1).

The reported results have implications both for social neuroscience and human-robot interaction: with regard to social neuroscience, this is the first report showing that mind perception is such a fundamental process when interacting with artificial agents that it cannot be suppressed even when it has negative effects on performance in conditions where a cognitive conflict is present. With regard to human-robot interaction, the experiments suggest that optimizing human appearance of robotic agents might not yield the desired outcome of improved performance but actually lead to decreased performance due to drainage of cognitive resources over time.

Acknowledgments

We would like to thank Jessica Langdorf, as well as our other research assistants for help with data collection.

Funding

This research did not receive any specific grant from funding agencies in the public, commercial, or not-for-profit sectors.

REFERENCES

- Abubshait, A., & Wiese, E. (2017). You Look Human, But Act Like a Machine: Agent Appearance and Behavior Modulate Different Aspects of Human–Robot Interaction. *Frontiers in psychology*, 8, 1393.
- Balas, B., & Tonsager, C. (2014). Face animacy is not all in the eyes: Evidence from contrast chimeras. *Perception*, 43(5), 355-367.
- Baron-Cohen, S., Leslie, A. M., & Frith, U. (1985). Does the autistic child have a theory of mind?. *Cognition*, 21(1), 37-46.
- Belopolsky, A. V., Devue, C., & Theeuwes, J. (2011). Angry faces hold the eyes. *Visual Cognition*, 19(1), 27-36.
- Bering, J. M., & Johnson, D. D. (2005). "O Lord... You perceive my thoughts from afar": Recursiveness and the evolution of supernatural agency. *Journal of Cognition and Culture*, 5(1), 118-142.
- Bornstein, R. F., & D'Agostino, P. R. (1992). Stimulus recognition and the mere exposure effect. *Journal of Personality and Social Psychology*, 63(4), 545-552.
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., & Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychological review*, 108(3), 624.
- Castelli, F., Happé, F., Frith, U., & Frith, C. (2000). Movement and mind: a functional imaging study of perception and interpretation of complex intentional movement patterns. *Neuroimage*, 12(3), 314-325.
- Carpenter, J. (2016). *Culture and Human-Robot Interaction in Militarized Spaces: A War Story*. Routledge.

- Cheetham M., Suter P., Jancke L. (2014). Perceptual discrimination difficulty and familiarity in the uncanny valley: more like a "Happy Valley." *Frontiers in Psychology*, 5(1219). 10.3389/fpsyg.2014.01219
- Cheetham, M., Suter, P., and Jancke L. (2011). The human likeness dimension of the "uncanny valley hypothesis": behavioral and functional MRI findings. *Frontiers in Human Neuroscience*. 5(126). 10.3389/fnhum.2011.00126
- Ciaramidaro, A., Adenzato, M., Enrici, I., Erk, S., Pia, L., Bara, B. G., & Walter, H. (2007). The intentional network: how the brain reads varieties of intentions. *Neuropsychologia*, 45(13), 3105-3113.
- Cikara, M., Bruneau, E. G., & Saxe, R. R. (2011). Us and them intergroup failures of empathy. *Current Directions in Psychological Science*, 20(3), 149-153.
- Craig, A. (1984). Human engineering: The control of vigilance. In J.S. Warm (Ed.), *Sustained Attention in Human Performance* (pp. 247-291). Chichester, England: Wiley.
- Davies, D. R., & Parasuraman, R. (1982). *The Psychology of Vigilance*. New York: Academic Press.
- Demoulin, S., Cortes, B. P., Viki, T.G., Rodriguez, A.P., Rodriguez, R.T., Paladino, M.P., & Leyens, J. P. (2009). The role of in-group identification in infra-humanization. *International Journal of Psychology*, 44 (1), 4-11.
- Dennett, D. (2003). Who's on first? Heterophenomenology explained. *Journal of Consciousness Studies*, 10(9-10), 19-30.
- Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual review of neuroscience*, 18(1), 193-222.

- Deska, J. C., Almaraz, S. M., & Hugenberg, K. (2016). Of mannequins and men: Ascriptions of mind in faces are bounded by perceptual and processing similarities to human faces. *Social Psychological and Personality Science*. 10.1177/1948550616671404.
- Egner, T., & Hirsch, J. (2005). Cognitive control mechanisms resolve conflict through cortical amplification of task-relevant information. *Nature neuroscience*, 8(12), 1784.
- Epley, N., Akalis, S., Waytz, A., & Cacioppo, J. T. (2008a). Creating social connection through inferential reproduction loneliness and perceived agency in gadgets, gods, and greyhounds. *Psychological Science*, 19(2), 114-120.
- Epley, N., Waytz, A., Akalis, S., & Cacioppo, J. T. (2008b). When we need a human: Motivational determinants of anthropomorphism. *Social Cognition*, 26(2), 143-155.
- Epley, N., Waytz, A., & Cacioppo, J. T. (2007). On seeing human: A three-factor theory of anthropomorphism. *Psychological Review*, 114(4), 864-886.
- Eysenck, M. W., & Calvo, M. G. (1992). Anxiety and performance: The processing efficiency theory. *Cognition & Emotion*, 6(6), 409-434.
- Eyssel, F., & Reich, N. (2013). Loneliness makes the heart grow fonder (of robots): On the effects of loneliness on psychological anthropomorphism. *Proceedings of the 8th ACM/IEEE international conference on Human-robot interaction*, 121-122. IEEE Press.
- Ferrey, A. E., Burleigh, T. J., & Fenske, M. J. (2015). Stimulus-category competition, inhibition, and affective devaluation: a novel account of the uncanny valley. *Frontiers in psychology*, 6, 249.

- Fox, E., Russo, R., Bowles, R., & Dutton, K. (2001). Do threatening stimuli draw or hold visual attention in subclinical anxiety?. *Journal of Experimental Psychology: General*, *130*(4), 681-700.
- Fox, E., Russo, R., & Dutton, K. (2002). Attentional bias for threat: Evidence for delayed disengagement from emotional faces. *Cognition & Emotion*, *16*(3), 355-379.
- Freeman, J. B., & Ambady, N. (2010). MouseTracker: Software for studying real-time mental processing using a computer mouse-tracking method. *Behavior Research Methods*, *42*(1), 226-241.
- Frith, C. D., & Frith, U. (2006). The neural basis of mentalizing. *Neuron*, *50*(4), 531-534.
- Friesen, C. K., Ristic, J., & Kingstone, A. (2004). Attentional effects of counterpredictive gaze and arrow cues. *Journal of Experimental Psychology: Human Perception and Performance*, *30*(2), 319.
- Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect perception of animacy irresistibly influences interactive behavior. *Psychological Science*, *21*(12), 1845-1853.
- Goldstone, R. L., & Hendrickson, A. T. (2010). Categorical perception. *Cognitive Science*, *1*(1), 69-78.
- Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology: General*, *121*(4), 480.
- Gray, K., Young, L., & Waytz, A. (2012). Mind perception is the essence of morality. *Psychological Inquiry*, *23*(2), 101-124.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, *315*(5812), 619.

- Gutsell, J. N., & Inzlicht, M. (2012). Intergroup differences in the sharing of emotive states: neural evidence of an empathy gap. *Social Cognitive and Affective Neuroscience*, 7(5), 596-603.
- Hackel, L. M., Looser, C. E., & Van Bavel, J. J. (2014). Group membership alters the threshold for mind perception: The role of social identity, collective identification, and intergroup threat. *Journal of Experimental Social Psychology*, 52, 15-23.
- Haley, K. J., & Fessler, D. M. (2005). Nobody's watching?: Subtle cues affect generosity in an anonymous economic game. *Evolution and Human Behavior*, 26(3), 245-256.
- Harnad, S. (1987) Psychophysical and cognitive aspects of categorical perception: A critical overview. In: S. Harnad (Ed.), *Categorical Perception: The Groundwork of Cognition*. New York: Cambridge University Press.
- Harris, L. T., & Fiske, S. T. (2011). Perceiving humanity or not: A social neuroscience approach to dehumanized perception. In: A. Todorov & T. Fiske (Eds.), *Social neuroscience: Toward understanding the underpinnings of the social mind*, pp. 123-134.
- Harris, L. T., & Fiske, S. T. (2006). Dehumanizing the lowest of the low neuroimaging responses to extreme out-groups. *Psychological Science*, 17(10), 847-853.
- Haslam, N. (2006). Dehumanization: An integrative review. *Personality and Social Psychology Review*, 10(3), 252-264.
- Hein, G., Silani, G., Preuschoff, K., Batson, C. D., & Singer, T. (2010). Neural responses to in-group and outgroup members' suffering predict individual differences in costly helping. *Neuron*, 68(1), 149-160.

- Hertz, N., & Wiese, E. (2016). Influence of Agent Type and Task Ambiguity on Conformity in Social Decision Making. In *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 60(1), 313-317.
- Kätsyri, J., Förger, K., Mäkärräinen, M., & Takala, T. (2015). A review of empirical evidence on different uncanny valley hypotheses: support for perceptual mismatch as one road to the valley of eeriness. *Frontiers in Psychology*, 6, 390.
- Lakatos, G., Gácsi, M., Konok, V., Brúder, I., Bereczky, B., Korondi, P., & Miklósi, Á. (2014). Emotion attribution to a non-humanoid robot in different social situations. *PloS One*, 9(12), e114207.
- Lee, A. Y. (2001). The mere exposure effect: An uncertainty reduction explanation revisited. *Personality and Social Psychology Bulletin*, 27(10), 1255-1266.
- Leyens, J. P., Paladino, P. M., Rodriguez-Torres, R., Vaes, J., Demoulin, S., Rodriguez-Perez, A., & Gaunt, R. (2000). The emotional side of prejudice: The attribution of secondary emotions to ingroups and outgroups. *Personality and Social Psychology Review*, 4(2), 186-197.
- Looser, C. E., Guntupalli, J. S., & Wheatley, T. (2013). Multivoxel patterns in face-sensitive temporal regions reveal an encoding schema based on detecting life in a face. *Social Cognitive and Affective Neuroscience*, 8(7), 799-805.
- Looser, C. E., & Wheatley, T. (2010). The tipping point of animacy: How, when, and where we perceive life in a face. *Psychological Science*, 21(12), 1854-1862.
- Lundqvist, D., Flykt, A., & Öhman, A. (1998). The Karolinska directed emotional faces (KDEF). CD ROM from Department of Clinical Neuroscience, Psychology section, Karolinska Institutet.

- Mackworth, N. H. (1948). The breakdown of vigilance during prolonged visual search. *Quarterly Journal of Experimental Psychology*, 1(1), 6-21.
- Manly, T., Hawkins, K., Evans, J., Woldt, K., & Robertson, I. H. (2002). Rehabilitation of executive function: Facilitation of effective goal management on complex tasks using periodic auditory alerts. *Neuropsychologia*, 40(3), 271-281.
- Martini, M. C., Gonzalez, C. A., & Wiese, E. (2016). Seeing minds in others—Can agents with robotic appearance have human-like preferences?. *PloS one*, 11(1), e0146310.
- Martini, M. C., Buzzell, G. A., & Wiese, E. (2015). Agent appearance modulates mind attribution and social attention in Human-Robot Interaction. *Proceedings of International Conference on Social Robotics*, 431-439. Springer International Publishing.
- Maurer, D., Le Grand, R., & Mondloch, C. J. (2002). The many faces of configural processing. *Trends in Cognitive Sciences*, 6(6), 255-260.
- Meng, M., & Tong, F. (2004). Can attention selectively bias bistable perception? Differences between binocular rivalry and ambiguous figures. *Journal of Vision*, 4(7), 2-2.
- Molloy, R., & Parasuraman, R. (1996). Monitoring an automated system for a single failure: Vigilance and task complexity effects. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 38(2), 311-322.
- Mori, M., MacDorman, K. F., & Kageki, N. (2012). The uncanny valley [from the field]. *Robotics & Automation Magazine*, 19(2), 98-100.
- Palermo, R., & Rhodes, G. (2007). Are you always on my mind? A review of how face perception and attention interact. *Neuropsychologia*, 45(1), 75-92.
- Parasuraman, R., de Visser, E., Clarke, E., McGarry, W. R., Hussey, E., Shaw, T., & Thompson, J. C. (2009). Detecting threat-related intentional actions of others: Effects of image

quality, response mode, and target cuing on vigilance. *Journal of Experimental Psychology: Applied*, 15(4), 275-290.

Parasuraman, R., Warm, J. S., & See, J. E. (1998). Brain systems of vigilance. In R. Parasuraman (Ed.), *The attentive brain* (pp. 221-256). Cambridge, MA, US: The MIT Press.

Parasuraman, R., & Mouloua, M. (1987). Interaction of signal discriminability and task type in vigilance decrement. *Perception & Psychophysics*, 41(1), 17-22.

Pollick, F. E. (2009). In search of the uncanny valley. *Proceedings of International Conference on User Centric Media*, 69-78. Springer Berlin Heidelberg.

Premack, D., & Woodruff, G. (1978). Does the chimpanzee have a theory of mind?. *Behavioral and Brain sciences*, 1(4), 515-526.

Repp, B. H. (1984). Categorical perception: Issues, methods, findings. In J. Lass (Ed.), *Speech and language: Advances in basic research and practice* (Vol. 10, pp. 243-335). New York: Academic Press.

Rosset, E. (2008). It's no accident: Our bias for intentional explanations. *Cognition*, 108(3), 771-780.

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., & Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science*, 300(5626), 1755-1758.

Schein, C., & Gray, K. (2015). The unifying moral dyad liberals and conservatives share the same harm-based moral template. *Personality and Social Psychology Bulletin*, 41(8), 1147-63.

- Shariff, A. F., & Norenzayan, A. (2007). God is watching you priming God concepts increases prosocial behavior in an anonymous economic game. *Psychological Science*, 18(9), 803-809.
- Shaw, T. H., Funke, M. E., Dillard, M., Funke, G. J., Warm, J. S., & Parasuraman, R. (2013). Event-related cerebral hemodynamics reveal target-specific resource allocation for both “go” and “no-go” response-based vigilance tasks. *Brain and cognition*, 82(3), 265-273.
- Singer, P. (2008). *Robots at War*. Wilson Quarterly.
- Smith, M. A., & Wiese, E. (2016). Look at me now: Investigating delayed disengagement for ambiguous human-robot stimuli. *Proceedings of International Conference on Social Robotics*, 950-960. Springer International Publishing.
- Sterzer, P., Kleinschmidt, A., & Rees, G. (2009). The neural bases of multistable perception. *Trends in cognitive sciences*, 13(7), 310-318.
- Temple, J. G., Warm, J. S., Dember, W. N., Jones, K. S., LaGrange, C. M., & Matthews, G. (2000). The effects of signal salience and caffeine on performance, workload and stress in an abbreviated vigilance task. *Human Factors*, 42, 183–194.
- Teufel, C., Fletcher, P. C., & Davis, G. (2010). Seeing other minds: Attributed mental states influence perception. *Trends in Cognitive Sciences*, 14(8), 376-382.
- Van Overwalle, F. (2009). Social cognition and the brain: A meta-analysis. *Human Brain Mapping*, 30(3), 829–858. <https://doi.org/10.1002/hbm.20547>
- Wagner, D. D., Kelley, W. M., & Heatherton, T. F. (2011). Individual differences in the spontaneous recruitment of brain regions supporting mental state understanding when viewing natural social scenes. *Cerebral Cortex*, 21(12), 2788-2796.

- Warm, J. S., Matthews, G., & Finomore, V. S. (2008). Vigilance, workload, and stress. *Performance Under Stress*, 115-141.
- Warm, J. S., Parasuraman, R., & Matthews, G. (2008). Vigilance requires hard mental work and is stressful. *Human Factors: The Journal of the Human Factors and Ergonomics Society*, 50(3), 433-441.
- Waytz, A., Gray, K., Epley, N., & Wegner, D. M. (2010). Causes and consequences of mind perception. *Trends in cognitive sciences*, 14(8), 383-388.
- Weis, P.P. & Wiese, E. (2017). Cognitive Conflict as Possible Origin of the Uncanny Valley. *Proceedings of the Human Factors and Ergonomics Society 2017 Annual Meeting*; pp. 1599-1603; Austin, TX.
- Wheatley, T., Weinberg, A., Looser, C., Moran, T., & Hajcak, G. (2011). Mind perception: Real but not artificial faces sustain neural activity beyond the N170/VPP. *PloS one*, 6(3), e17960.
- Wickens, C. D., & Kessel, C. (1980). Processing resource demands of failure detection in dynamic systems. *Journal of Experimental Psychology: Human Perception and Performance*, 6(3), 564-577.
- Wiese, E., Wykowska, A., & Müller, H. J. (2014). What we observe is biased by what other people tell us: Beliefs about the reliability of gaze behavior modulate attentional orienting to gaze cues. *PloS one*, 9(4), e94529.
- Wiese, E., Wykowska, A., & Müller, H. J. (2013). Making eyes with robots: Readiness to engage in Human-Robot-interaction depends on the attribution of intentionality. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*, 57(1), 1174-1178. SAGE Publications.

Wiese, E., Wykowska, A., Zwickel, J., & Müller, H. J. (2012). I see what you mean: how attentional selection is shaped by ascribing intentions to others. *PloS one*, 7(9), e45391.

Wittgenstein, L. (1953). *Philosophical investigations*. Oxford: Blackwell.

Wykowska, A., Wiese, E., Prosser, A., & Müller, H. J. (2014). Beliefs about the minds of others influence how we process sensory information. *PLoS one*, 9(4), e94339.

Zlotowski, J., Proudfoot, D., & Bartneck, C. (2013). More human than human: Does the uncanny curve really matter? *Proceedings of the HRI2013 Workshop on Design of Humanlikeness in HRI*, 7-13.

Figure 1. Agent stimuli arranged by degree of physical humanness with the least human on the left: robot (0% human), robotoid (30% human), humanoid (70% human), and human (100% human).

Figure 2. Probes in the vigilance task: Gun (*left*) and hair dryer (*right*). Stimuli were selected for their morphological similarity but discriminant threat-relevance. During each trial, one of the probes appeared on the right or left side of the screen and was either looked-at or not looked-at by the stimulus. Participants responded only on critical trials, that is when the agent looked towards the gun.

Figure 6. Average change of hit rates in each block by agent type in (a) Experiment 1, (b) Experiment 2, and (c) Experiment 3. Experiment 1 showed a significant vigilance decrement across conditions, with a significantly stronger decrement for the humanoid compared to the human and robot stimulus. Experiment 2 also showed a significant vigilance decrement across conditions, but there was no significant difference in vigilance decrement between agents. Experiment 3 did not show a vigilance decrement. Error bars represent standard errors.

Figure 3. Trial sequence in the vigilance task: A fixation cross is presented for 250 ms, followed by an agent's face looking straight ahead for 500 ms before changing gaze direction either looking to the left or right side of the screen. 200 ms later, the probe appears either at the gazed-at location or opposite of where the face is looking for 100 ms. Afterwards, a blank screen is presented for 1200 ms and participants are supposed to give their response.

Figure 4. Results of Experiment 1: Average rate of change in hit rates, false alarm rates and reaction times in each block by agent type. (a) *Hit rates*: there was a significant decrease in hit rates over time, with a significantly stronger decrease in the humanoid condition compared to all other conditions. (b) *False alarm rates*: there was an increase in false alarm rate over time and higher false alarm rates for the humanoid condition compared to all other conditions, but no significant difference in vigilance performance between agents. (c) *Reaction times*: there was a significant overall increase in reaction times over time, but no significant difference in vigilance performance between agents (error bars represent standard errors).

Figure 5. Mind ratings and categorization data to test for cognitive conflict: (a) Participants evaluated the stimuli used in the current study in terms of their capability of having a mind (on a 7-point Likert scale where 1= Definitely Does Not Have a Mind and 7= Definitely Has a Mind): the human agent was clearly rated as having a mind (6.20 out of 7), while the robotoid (1.76 out of 7) and robot (1.63 out of 7) agents were clearly rated as not having a mind. The humanoid agent received intermediate ratings of mind perception (3.95 out of 7). (b) When participants were asked to categorize the agents into either the human or non-human category in a forced-choice paradigm, there was much less consensus over the humanoid agent (42.2% human categorizations) compared to the human agent (99.4% human categorizations) and the robot and robotoid agents (1.9 % and 2.5% human categorizations, respectively); (c) furthermore, participants took significantly longer for the categorization of the humanoid agent (3.67 seconds) compared to the human, robot and robotoid agents (2.22 and 2.47 seconds, respectively).

Table 1. Average Hits Rates (%) across agents (Standard Error in parentheses)

Condition	Block							
	1	2	3	4	5	6	7	8
Robot	97.8 (1.2)	99.3 (0.7)	97.8 (1.2)	97.8 (1.6)	91.9 (3.5)	93.3 (3.7)	92.6 (4.8)	94.8 (4.5)
Robotoid	100.0 (0)	98.4 (1.1)	96.8 (2.5)	96.0 (1.9)	96.8 (1.4)	93.7 (2.2)	95.2 (3.2)	96.8 (1.8)
Humanoid	100.0 (0)	95.2 (2.8)	92.9 (2.8)	90.5 (3.7)	85.7 (5.4)	80.2 (5.8)	81.7 (5.8)	72.2 (8.2)
Human	93.3 (2.8)	95.6 (3.7)	94.1 (2.1)	92.6 (2.5)	92.6 (3.9)	93.3 (3.6)	94.8 (2.6)	97.8 (1.2)

Table 2. Average False Alarms (%) for Each Agent (Standard Error in parentheses).

Condition	Block							
	1	2	3	4	5	6	7	8
<i>Robot</i>	3.0 (0.8)	2.3 (0.6)	1.5 (0.4)	1.6 (0.5)	1.0 (0.3)	0.6 (0.4)	0.8 (0.2)	1.1 (0.4)
<i>Robotoid</i>	2.5 (0.5)	1.5 (0.5)	0.8 (0.3)	0.3 (0.2)	0.7 (0.4)	0.6 (0.2)	0.3 (0.2)	0.2 (0.2)
<i>Humanoid</i>	5.3 (2.4)	1.8 (0.5)	2.5 (0.8)	2.6 (0.8)	4.0 (1.1)	2.3 (0.7)	2.9 (0.9)	2.7 (1.1)
<i>Human</i>	4.7 (1.2)	4.1 (1.3)	2.8 (0.8)	2.1 (0.9)	1.5 (0.5)	2.1 (0.8)	3.1 (0.9)	1.4 (0.4)

Table 3. Average Reaction Times (in seconds) for Each Agent (Standard Error in parentheses).

Condition	Block							
	1	2	3	4	5	6	7	8
<i>Robot</i>	461.2 (14.0)	486.8 (19.0)	494.6 (20.5)	499.2 (20.1)	519.4 (21.3)	501.1 (23.1)	508.9 (26.4)	535.1 (32.3)
<i>Robotoid</i>	515.1 (18.9)	565.9 (26.6)	590.2 (22.3)	582.5 (21.9)	598.2 (22.7)	597.6 (27.9)	573.2 (23.0)	589.2 (27.1)
<i>Humanoid</i>	496.0 (26.6)	520.6 (27.9)	526.4 (27.9)	543.5 (30.1)	572.9 (32.0)	549.4 (25.4)	594.2 (29.4)	555.2 (40.0)
<i>Human</i>	462.9 (20.2)	514.9 (22.2)	554.1 (28.2)	532.8 (20.1)	558.8 (30.4)	537.1 (22.3)	535.7 (29.8)	566.2 (29.8)

Table 4. Social Judgment Task answered on a 7-point Likert scale (1= Definitely Not; 7= Definitely).

- 1) Rate how much this face looks alive.
- 2) Rate how much this face looks like it has a mind.
- 3) Do you think this agent would feel pain if it tripped and fell on hard ground?
- 4) Do you think this agent likes to hang out with friends?
- 5) Do you think this agent would be an interesting conversationalist?

Table 5. Average Hits Rates (%) for Each Agent in Experiment 2 (Standard Error in parentheses).

Condition	Block							
	1	2	3	4	5	6	7	8
<i>Robot</i>	94.8 (1.8)	98.5 (1.0)	96.2 (2.1)	94.8 (1.8)	91.9 (3.2)	88.1 (3.8)	94.8 (2.1)	93.3 (2.6)
<i>Humanoid</i>	98.5 (1.0)	99.3 (0.7)	99.3 (0.7)	97.8 (1.2)	94.8 (2.6)	94.8 (2.6)	94.8 (1.5)	91.9 (4.4)
<i>Human</i>	99.2 (0.8)	97.6 (1.3)	95.2 (2.8)	96.0 (2.5)	98.4 (1.6)	93.7 (3.3)	96.8 (1.4)	98.4 (1.1)

Table 6. Perceptual judgment task: Questions and response items.

- 1) Rate how old you think the face is.
- 2) Rate how attractive you think the face is.
- 3) Rate which fruit shape the face resembles.
- 4) Rate the gender appearance of the face.
- 5) Rate how long the hair of the agent is.

Table 7. Average Hits Rates (%) for Each Agent in Experiment 3 (Standard Error in parentheses).

Condition	Block							
	1	2	3	4	5	6	7	8
<i>Robot</i>	99.2 (0.8)	95.2 (1.9)	96.8 (1.4)	94.4 (2.3)	98.4 (1.1)	97.6 (1.7)	96.8 (1.8)	92.9 (3.0)
<i>Humanoid</i>	99.2 (0.7)	97.0 (1.7)	94.8 (2.4)	94.1 (2.8)	94.1 (2.8)	95.6 (2.4)	95.6 (2.4)	94.1 (3.2)
<i>Human</i>	97.8 (1.6)	97.0 (1.7)	98.5 (1.0)	99.3 (0.7)	98.5 (1.0)	99.3 (0.7)	98.5 (1.5)	99.3 (0.7)