


Cognitive Psychology

Faces Do Not Attract More Attention Than Non-Social Distractors in the Stroop Task

Anna Henschel¹ , Hannah Bargel², Emily S. Cross³  ^a

¹ Institute of Neuroscience and Psychology, University of Glasgow, Glasgow, Scotland, UK, ² Universität zu Lübeck, Lübeck, Germany, ³ Institute of Neuroscience and Psychology, University of Glasgow, Glasgow, Scotland, UK; Department of Cognitive Science, Macquarie University, Sydney, Australia

Keywords: social motivation, eye contact effect, salient social cues, robots, pareidolia

<https://doi.org/10.1525/collabra.18521>

Collabra: Psychology

Vol. 7, Issue 1, 2021

As robots begin to receive citizenship, are treated as beloved pets, and given a place at Japanese family tables, it is becoming clear that these machines are taking on increasingly social roles. While human-robot interaction research relies heavily on self-report measures for assessing people's perception of robots, a distinct lack of robust cognitive and behavioural measures to gauge the scope and limits of social motivation towards artificial agents exists. Here we adapted Conty and colleagues' (2010) social version of the classic Stroop paradigm, in which we showed four kinds of distractor images above incongruent and neutral words: human faces, robot faces, object faces (for example, a cloud with facial features) and flowers (control). We predicted that social stimuli, like human faces, would be extremely salient and draw attention away from the to-be-processed words. A repeated-measures ANOVA indicated that the task worked (the Stroop effect was observed), and a distractor-dependent enhancement of Stroop interference emerged. Planned contrasts indicated that specifically human faces presented above incongruent words significantly slowed participants' reaction times. To investigate this small effect further, we conducted a second experiment (N=51) with a larger stimulus set. While the main effect of the incongruent condition slowing down participants' reaction time replicated, we did not observe an interaction effect of the social distractors (human faces) drawing more attention than the other distractor types. We question the suitability of this task as a robust measure for social motivation and discuss our findings in the light of recent conflicting results in the social attentional capture literature.

Introduction

Glancing upon Giuseppe Arcimboldo's famous 16th century artwork "Air", a collection of colourful birds transforms into the side profile of an elegant man. The effect Arcimboldo cleverly applied to many of his paintings is known as *pareidolia*, which describes the illusory perception of human faces in random patterns. This tendency is not only capitalized on in the arts, online communication, and product design, but also in research, where variations on the visual illusion are used to investigate mechanisms of face perception (Bubic et al., 2014; Guido et al., 2019; Martinez-Conde et al., 2015; Pavlova et al., 2018; Robertson et al., 2017; Wodehouse et al., 2018).

While the origin of the pareidolia phenomenon is somewhat contentious (with explanations ranging from "visual false alarms" to reflecting a deeply ingrained need for social contact), it points to the fact that human faces have a unique status in our visual environment (DiSalvo & Gem-

perle, 2003; Wodehouse et al., 2018; Zhou & Meng, 2019). From birth, babies exhibit a preference for gazing at faces compared to scrambled faces, with a bias for gazing at others' eyes developing within the first year of life (Hessels, 2020). Replications of a seminal eye-tracking study by Yarbus (1967) confirm that participants invariably have a gaze preference for people, faces and eyes (DeAngelus & Pelz, 2009). Faces are a rich source of information, giving insight into another person's emotions, their intentions, and their personality traits. Willis and Todorov (2006), for example, have shown that the proverb "you only get one chance to make a first impression" is grounded in empirical truth. They found that participants were able to make reliable trait judgements on attractiveness, likeability, trustworthiness, competence and aggressiveness within split seconds. In yet another study, perceivers were capable of deducing the social class of unfamiliar faces above chance level, highlighting the importance of face perception and its potential societal impact (Bjornsdottir & Rule, 2017).

^a Corresponding author: Emily.cross@mq.edu.au

An integrative theoretical account on the relative importance of social cues, such as faces, by Chevallier and colleagues describes social motivation by means of three main components: social reward, social maintaining, and social orienting (2012). Interactions with others, the authors argue, are inherently rewarding, relationships are driven by our goals to maintain and improve them, and social cues are thus prioritized. The authors propose that social motivation is determined by specialized biological processes, which developed due to an evolutionary advantage of collaborating with other humans. Thus, social information in the form of facial cues is thought to be extremely powerful in terms of claiming attentional resources, increasing our chances for improved coordination and cooperative work with others (2012).

Given their prioritization in our visual environment, it is unsurprising that faces have been the central focus of many visual attention studies. Collectively, these studies point towards faces ranking above objects in capturing automatic attention. Using a change blindness paradigm, Ro, Russel and Lavie (2001) found that participants detected changes in temporarily presented faces more quickly than changes in any other object. This effect disappeared when the face stimuli were inverted. Automatic attentional capture by faces was further investigated by Theeuwes and Van der Stigchel (2006), who criticized that Ro and colleagues' (2001) results could have been due to merely a preference for attending to faces, and not reflective of truly exogenous attentional capture. In their inhibition of return paradigm, these authors found evidence for automatic attentional capture induced by faces as compared to object stimuli. The authors observed a delayed gaze response towards locations that had previously shown a face and reasoned that this represented true attentional capture by faces, rather than difficulties with disengaging attention from them. Bindemann and colleagues (2007) sought to understand whether attentional capture by facial cues could be entirely determined by their salience, or whether this effect is also modified endogenously, by participants' own volition. As a matter of fact, participants were able to direct their attention away from faces towards objects when these were more predictive of the cued target location in a dot-probe paradigm. However, the authors claimed an overall face bias persisted, with participants showing greater ease at directing attention to predictive faces versus predictive objects. Experiments by Langton and colleagues (2008) further affirmed the notion that attentional capture by faces is automatic and involuntary. Searching a visual array for a butterfly was slowed by the presence of an "additional singleton", a task-irrelevant face. Here, the authors concluded that humans became consciously aware of faces before any other none-face item. Overall, a large body of evidence suggests that social attentional capture by facial cues is a robust phenomenon, providing evidence for the putative social orienting pillar of the social motivation model.

Beyond seeing faces in oddly shaped clouds, Martian craters or pieces of burnt toast, we also encounter deliberate pareidolic design when we interact with humanoid robots (DiSalvo et al., 2002; DiSalvo & Gemperle, 2003; Wodehouse et al., 2018). Due to the face's role in communicating emotions, and more generally, facilitating social

interactions, the design of human-like (or at least human-readable) robot faces has attracted considerable attention and investment in the domain of social robotics. A key driver behind humanoid robot design is the desire to build a believable social agent, while mitigating the potential damaging effects an overly human-like appearance could have on the user (e.g., coming too close to the so-called "uncanny valley"; DiSalvo & Gemperle, 2003). Thus, in order to avoid an uncanny experience, or over-promise on the robot's functionality, a popular design choice for socially assistive robots is a humanoid face with simple geometric shapes alluding to familiar, human features (Kalegina et al., 2018). Indeed, when participants were asked to rate the humanness of humanoid robot heads, only a few features accounted for more than 62% of variance: the eyes, eyelids, nose and mouth (DiSalvo et al., 2002). This is in line with a study by Omer and colleagues, which mapped the features that contributed to the global gestalt of pareidolia faces, identifying the eyes and the mouth (2019). Robots' facial cues are viewed as one of the crucial four dimensions in driving human-likeness ratings, and in a survey of humanoid robots, 87.5% had at least some facial features (DiSalvo et al., 2002; Phillips et al., 2018). It is of note that when establishing an impression of animacy, viewing the face as a whole is crucial, with participants being more hesitant to make judgements about the presence of mind in an agent when viewing cropped facial cues in isolation (Looser & Wheatley, 2010). Hence, and as Geiger and Balas (2020) point out, robot faces, which we have presented here as a special case of intentional pareidolia, constitute a border category of face processing, and while some research exists on attentional capture by pareidolic faces, less is known about the social relevance of robot faces. This question however is crucial, as humanoid robots become increasingly commonplace in modern society, taking on care, companionship and support roles. Hence, an important goal is to develop robust behavioural tasks that probe the relevance of robotic, compared to human, social cues.

Research on pareidolic faces and the extent to which they engage social attentional processes has yielded mixed results so far, with some researchers arguing for the crucial role of top-down information driving the face illusion effect (Takahashi & Watanabe, 2013, 2015), and others providing evidence for a bottom-up account of the phenomenon (Liu et al., 2014; Robertson et al., 2017). Takahashi and Watanabe (2013) investigated reflexive attentional shifts induced by pareidolic faces using a gaze cueing paradigm. The authors found a cueing effect of pareidolic faces, however, this effect disappeared when participants were not explicitly instructed that the presented objects could be interpreted as faces. In a follow-up study, Takahashi and Watanabe (2015) found that face awareness, i.e. perceiving an object (here: three dots arranged as a triangle) as a face improved participants performance on a target detection task. This advantage disappeared when subjects were instructed to detect a triangle target shape, rather than a face target. The authors concluded that despite their identical shape, faces receive prioritized further processing due to top-down modulation of face awareness. On the other hand, a study by Ariga and Arihara (2017) did not find that pareidolia faces captured visual attention when presented as task-irrelevant distractors

in a letter identification task. However, when human faces were presented as distractors among a rapid serial presentation of letters, accuracy was significantly impaired. There was no difference between pareidolia faces and their defocused control images for any of the various time lag conditions in the letter identification task. While Ariga and Arihara (2017) conclude that attentional capture by facial cues is exclusively reserved for human faces, yet another study shows that pareidolia faces were able to elicit deeper forms of social engagement, surpassing an initial face detection stage and eliciting further specialized processing. In their study, Palmer and Clifford (2020) presented pareidolic stimuli exhibiting directional eye gaze and found that during a subsequent human direct eye gaze task, sensory adaptation had taken place: the illusory faces influenced the perception of the human face stimuli. This finding is at odds with Robertson, Jenkins and Burton's (2017) conclusion: these authors claim that their participants' performance on several pareidolia face detection tasks was unrelated to their performance on face identification tasks, suggesting a functional dissociation and no higher-level face processing taking place elicited by illusory faces.

While the evidence on how deeply illusory faces are perceived as social is mixed, they constitute an ideal control for human facial features in social attentional capture tasks. This also raises the question how deliberate pareidolic faces, such as humanoid robots, might engage our visual attention, as these agents are capable of at least some interactions with the physical world. Some preliminary evidence exists from an electrophysiological study by Geiger and Balas (2020), which suggests that robot faces were more likely to be perceived as objects, rather than faces when presented in an inversion effect paradigm. The authors found that the face sensitive N170 ERP-component was moderately influenced by robot faces, ranking somewhere between objects such as clocks and real or computer-generated human faces.

The neuronal architecture underlying the prioritization of social cues has been shown to include both cortical and subcortical regions, including the amygdala, the ventral striatum, the orbitofrontal cortex and the ventromedial prefrontal cortex. These brain structures, which are reliably engaged during other types of reward processing as well, seem to be sensitive to, or perhaps even signal, the importance of social aspects of our environment (Schilbach et al., 2011). A formal theory in favour of a specialized subcortical fast track was put forward by Senju and Johnson, who coined the "eye contact effect" (2009). The fast-track modulator model claims that eye contact receives prioritized processing via a subcortical route. To test this hypothesis, Conty and colleagues (2010) conducted experiments on the distracting effect of social cues while participants were engaged in a cognitively demanding task: the classic colour Stroop paradigm (MacLeod & MacDonald, 2000; Stroop R, 1935).

Despite the above reviewed variety of paradigms which probe (social) attentional capture, the Stroop task has proven to be a particularly popular vehicle. Named after the psychologist who discovered the effect, hundreds of studies have shown that naming the ink colour of an incongruent colour word (i.e., the word "RED" presented in green) produces slower reaction times than determining the colour

of a control word (the letters "XXX" presented in green). This interference effect, which highlights the fact that task-irrelevant information is processed concomitantly and automatically, has inspired a multitude of extensions, including pictorial, spatial, and social versions (MacLeod & MacDonald, 2000). For example, in the facial-emotional Stroop, participants name the ink colour of emotional, compared to neutral faces, which are overlaid with a coloured filter. Past research has shown that sad participants and participants with higher trait anger are slower to name the colour of angry versus neutral faces (Isaac et al., 2012; Van Honk et al., 2000; van Honk et al., 2001). Thus, the Stroop task has been validated as a suitable paradigm to assess the distracting power of task-irrelevant information, such as facial cues.

In Conty and colleagues' study (2010), the cropped eye-regions of human faces with open or closed eyes - in one of two head orientations - were presented as task-irrelevant distractors on top of the Stroop task. The authors found that the interference effect produced by the competition between the automatic processing of word meaning and ink colour was further enhanced in the *direct gaze* condition, regardless of the head orientation. In a follow-up experiment, Conty et al. (2010) showed participants visual gratings and grey colour blocks as distractors, which the authors argue excluded the possibility that the effect might have been driven by low-level visual properties of the images - as open eyes have an inherently stronger visual contrast than closed eyes. In a third experiment with a new participant sample, they again found no difference between closed or averted eyes when presented as distractors on the task. Conty and colleagues conclude that the salience of direct eye contact was so strong that it tapped into processing resources needed to perform well on the main task: responding quickly and accurately to the target words (2010).

A later study from the same lab by Chevallier and colleagues replicated and extended the costly eye contact effect (2013). Importantly, the authors tested the paradigm in two groups of children: typically developing boys and a group of male adolescents with Autism Spectrum Condition (ASC). Again, open and closed eyes were presented as distractors above the neutral and incongruent words, however, this time a non-social control condition was added: flower images. As expected, the authors report the Stroop interference effect, where incongruent words significantly slowed participants' reaction times. The typically developing group showed the hypothesized enhanced interference in the social condition (here open and closed eyes were taken together as the 'social' category), while the ASC group showed the opposite effect. However, when investigating only the open versus closed eyes, stronger interference for open eyes was preserved in adolescents with ASC. The authors interpreted their findings as yet another confirmation for the strong salience of task-irrelevant social distractors but remark that their results are limited by their specific stimulus set and invite future studies to investigate other types of social distractors, such as whole faces.

In the current study, we built on their paradigm by testing the extent to which human, robot or object faces capture attention automatically, by presenting them on top of the classic colour Stroop task. We were interested in extending the Stroop paradigm to test a wider variety of social cues in

terms of their motivational value, as well as in evaluating the utility of the social Stroop task with robot faces as a valid behavioural task to probe social perception in HRI research.

Hypotheses. In line with a large body of literature on the Stroop interference effect, we expected that incongruent words would slow reaction times in comparison to the neutral target word condition, leading to the classic interference effect (MacLeod & MacDonald, 2000). Based on the findings by Conty et al. (2010) and Chevallier et al. (2013), as well as the established literature on social attentional capture, we further predicted that the more socially salient a cue is, the more it would lead to enhanced Stroop interference in this conceptual extension of the paradigm. The most socially salient stimuli used in the present study were human faces, which we predicted would increase reaction times in the incongruent Stroop condition. Less salient distractors were the robotic faces, which in theory allow for a more minimal form of social interaction. Even less socially salient distractors, the object (pareidolic) faces, contained facial cues but no capacity for the object to interact with the world in a social manner. Finally, we expected the control images, which held no social relevance whatsoever, to have no effect on reaction times in the incongruent condition of the Stroop task.

Experiment 1

Method

Preregistration and data statement. The experiment was pre-registered via www.AsPredicted.org. The document can be found at <https://osf.io/ky4b7/>. We report all measures in the experiment, all manipulations, any data exclusions and the sample size determination rule (Simmons et al., 2012). Data and the R analysis scripts are available (<https://osf.io/xyz4m/>). Due to copyright restrictions, the full stimulus set is not openly available, however it can be shared upon request.

Participants. An a-priori power analysis based on the contrast of interest resulted in a total sample size of 47 participants ($d_z=0.49$, $\alpha=0.05$, power=0.95, noncentrality parameter = 3.359, critical $t=1.678$, $Df=46$, actual power=0.95). We recruited 50 participants, however, based on our pre-registered exclusion criteria (diagnosis of ASD and having had a previous interaction with a robot) we excluded 9 participants. Two additional participants had insufficient English language skills, and thus the total number of exclusions was 11. The pre-registered exclusions were made based on participant answers on the experiment questionnaires' self-report items (for example: "Do you have a diagnosis of Autism Spectrum Disorder?" and "Have you interacted with a robot before?"). The other exclusions had to be made in addition, based on the difficulties of the participants with the task. We report a final sample size of $N=39$. Of the 39 participants, 26 were female, and reported a mean age of 27.41 years ($SD=7.35$). Ethical approval was obtained from the University of Glasgow ethics review board (300170224). All participants provided written informed consent prior to taking part and were reimbursed for their participation by payment. As in the original study, the experiment was framed as an experiment on colour perception.




Figure 1. A representation of the four different stimulus categories: human faces, robot faces, pareidolic faces and the control images, flowers. The human, robot and object distractors all have a direct gaze orientation and show a neutral facial expression. The full stimulus set is available upon request, as individual images are restricted by copyright.

Stimuli. A new stimulus set was built for this adapted version of the Stroop paradigm (Figure 1). The human faces were selected from neutral, frontally oriented facial expressions in the Radboud Faces Dataset and the London Faces Database (DeBruine & Jones, 2017; Langner et al., 2010). The robot and object faces, as well as the flowers, were selected from Google, with the aim to include only neutral, frontally-oriented faces. The rationale behind including only neutral faces was that emotional facial cues have been shown to draw attention, especially in comparison to neutral facial expressions (Pessoa et al., 2002; Theeuwes & Van der Stigchel, 2006; Vuilleumier, 2002).

An independent sample rated the first pool of human and robot images, resulting in a pre-selection of more neutrally perceived faces (more details can be found in the [Supplementary Materials](#)). Twelve unique images were obtained in each of the 4 categories and were edited to achieve a standard round form, mirrored, transformed to grey-scale and averaged according to mean contrast and luminance using the SHINE toolbox in MATLAB (Willenbockel et al., 2010). This resulted in 96 unique images in Experiment 1 (i.e. 24 per each of the four distractor conditions). Since the overall number of trials was 192 (closely modelled on the original study by Conty, Gimmig, et al., 2010), the distractor images were presented twice.

Procedure. Participants were tested in a quiet, dark cubicle on a computer, sitting 50 cm away from the screen. Participants familiarized themselves with the key responses in two training rounds. In the first training, colour-unrelated words (such as "BOWL" or "HAT") were presented in red, yellow, blue and green ink. Words low in arousal and with a medium valence score from the Affective Norms for English Words (Bradley & Lang, 1999) were selected. In this first practice block, participants received feedback on their performance accuracy and speed, whereas in the second round, the feedback was removed. Each practice block consisted of 48 trials. The experiment was split in 4 blocks, with short breaks after 48 trials. In total, the experiment took 25 minutes to complete.

Table 1. Mean reaction times and standard errors in milliseconds (Experiment 1).

		Humans	Robots	Objects	Flowers
Incongruent target	<i>M(SE)</i>	843 ± 11	807 ± 11	815 ± 11	796 ± 11
Neutral target	<i>M(SE)</i>	753 ± 10	768 ± 11	763 ± 10	760 ± 10

An experimental trial consisted of a centrally presented fixation cross, whose duration was jittered between 800 and 1300 milliseconds (Figure 2). After the fixation cross, the target word appeared, which extended horizontally over 1° of visual angle, and vertically over 0.5° of visual angle. Directly above the target words, the distractors were presented, extending over ca. 6° of visual angle. The images and word pairs remained on the screen until a response was made. There were equal numbers of incongruent and neutral Stroop trials, and no restrictions regarding the switch between incongruent and neutral trials were put in place, as they were presented randomly. The target word and distractor image pairs were fixed. Due to an error when setting up the PsychoPy experiment (Peirce, 2007), only female human faces were presented in the incongruent condition of the Stroop task, with all the male faces presented in the neutral condition. The object and robot distractor images in Experiment 1 were not one-to-one controlled by their mirror images across the incongruent and neutral conditions.

Statistical analysis (pre-registered). The percentage of accurate responses was calculated and analysed by means of a repeated measures ANOVA. For the analysis of the reaction times, incorrect responses were excluded, as were RTs that were two standard deviations above the mean or below 200ms. As a result, 606 trials (8.09%) were discarded (a detailed breakdown of the trial number per condition can be found in the [Supplementary Materials](#)).

We calculated a two-way repeated measures ANOVA with the target (incongruent vs. neutral) and distractor (human, robot, object, flower) as within-subjects conditions. Finally, we conducted planned contrasts. All analyses were conducted in R 3.5.3 (2019), using the {ez}, {psych}, {afex} and {emmeans} packages (Lawrence, 2016; Lenth et al., 2019; Revelle, 2018; Singmann et al., 2019).

Results

Accuracy. The repeated measures ANOVA showed a main effect of target, suggesting participants were more accurate in the neutral target word condition: $F(1, 38) = 7.48, p = 0.009, \eta^2 = .03$. However, the overall accuracy was very high (95.72%) and the effect size is considered small, so this was not further investigated.

Reaction times. A second repeated measures ANOVA was

calculated and as predicted, we saw a main effect of target, with incongruent words slowing down the reaction times of the participants: $F(1, 38) = 39.24, p < .001, \eta^2 = .03$. This finding confirms that our modified task was still effective at inducing a Stroop interference effect. In addition, we observed a small interaction effect of target x distractor: $F(3, 114) = 2.69, p = .049, \eta^2 = .003$. To investigate the difference in reaction times between specific conditions (comparing the effect of the human distractors in the incongruent condition with the flower distractors in the incongruent condition), planned contrasts were computed.

They revealed that the human faces were significantly more distracting than the flower images in the incongruent condition: $t(227) = -2.95, p = .004$ and drew more attention than the robotic faces as well ($t(227) = -2.15, p = .03$), but there was no significant difference to the object faces: $t(227) = -1.86, p = .06$. The Stroop interference scores (neutral trials subtracted from incongruent trials) are visualized in [Figure 3](#) and the mean reaction times with standard errors are summarized in [Table 1](#).

Discussion

In Experiment 1 we found an interaction effect in the predicted direction: human faces drew more automatic attention than flower images and robot faces, leading to enhanced interference in the Stroop task. However, the interaction that emerged, as evaluated by the ANOVA, was very small and just above the set significance level ($p = .049$). In addition, due to our conservative participant exclusion criteria, we experienced a larger drop-off in overall subject number than expected. Thus, the experiment was perhaps not adequately powered to detect the effect of interest. Furthermore, we speculated that the effect may have been influenced by the repetition of the distractor images, or due to the described programming error. We next decided to run the same paradigm again, this time recruiting a sufficiently large subject number (accounting for a drop-out rate of approximately 15-20% of participants), presenting both male and female faces in the incongruent Stroop condition, and doubling the number of unique distractors, thus preventing repeated viewing of the stimuli.

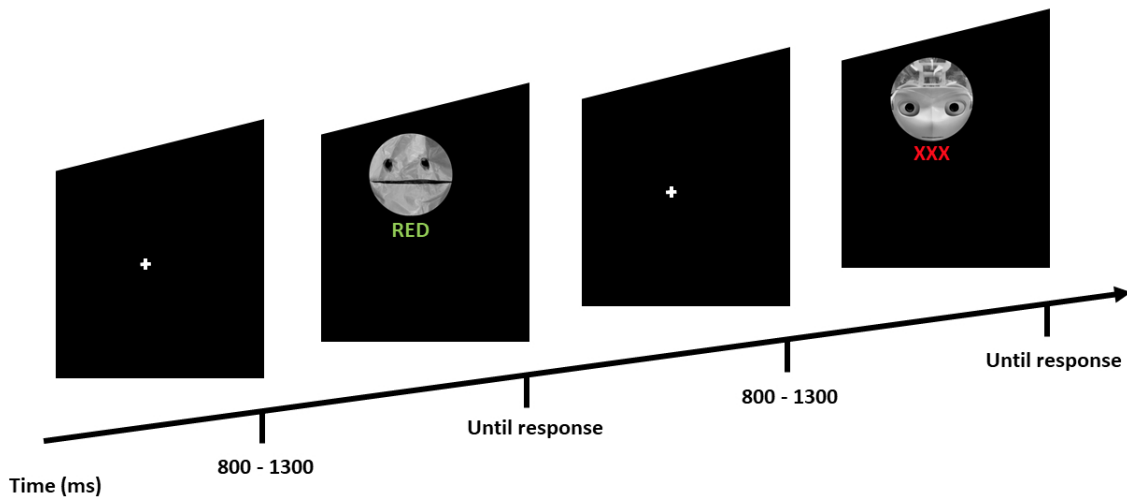


Figure 2. Schematic representation of a trial time course.

Experiment 2

Method

Preregistration and data statement. We followed the same procedures that were described in our preregistration document, as reported in Experiment 1.

Participants. A new set of participants ($N=70$) was recruited. In addition to the pre-registered exclusion criteria (outlined in Experiment 1 - Method), we added the condition of not having participated in the first experiment. After subject exclusion, 51 participants remained in the sample (39 female). The participants' mean age was 23.24 years ($SD=6.27$). All participants provided written informed consent prior to volunteering for this experiment and were reimbursed by payment. The experiment was approved by the University of Glasgow ethics review board (300180052).

Stimuli. The stimulus set was extended to include 12 new unique images for each distractor condition, which were mirrored and edited in the same way as outlined in Experiment 1. In total, we now had 192 unique distractors.

Procedure. The same experimental procedure was followed as described in Experiment 1. Following the completion of the Stroop task, we also asked participants to rate the unique (unmirrored) distractors based on agency (ability to plan and act) and experience (ability to sense and feel), to establish that the distractor categories were indeed perceived differently, with regard to their varying levels of social saliency. Participants rated each of the 96 images on both characteristics using a sliding scale from 0 to 100 in FormR (Arslan et al., 2019). The questions were derived from Gray, Gray and Wegner's study (2007) on mind perception of different kinds of agents. We used mind perception as a socialness proxy to distinguish between the control condition (flowers), inanimate (robot and pareidolic faces) and agents with a mind (humans). The analysis of the rat-

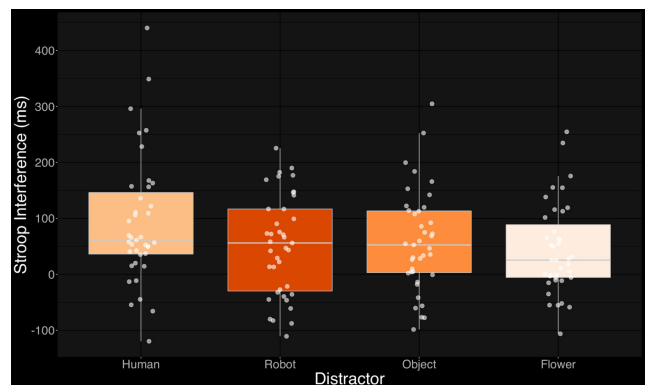


Figure 3. The Stroop interference scores were calculated by subtracting the neutral from the incongruent trials. Here the mean Stroop interference scores are shown for each of the distractor categories in Experiment 1.

ings confirmed that the stimulus categories were perceived differently: the human images received the highest agency and experience ratings. A detailed report of the stimulus ratings can be found in the [Supplementary Materials](#).

Statistical analysis. We followed the same data cleaning and analysis procedure as in Experiment 1. Incorrect trials were excluded, as well as reaction times below 200ms or 2 standard deviations above the mean (i.e. 1910ms). With this reaction time trimming criterion, we discarded 1061 trials (10.84%). A detailed breakdown of the number of trials remaining per condition can be found in the [Supplementary Materials](#).

Results

Accuracy. The repeated measures ANOVA showed no sig-

Table 2. Mean reaction times and standard errors in milliseconds (Experiment 2)

		Humans	Robots	Objects	Flowers
Incongruent target	<i>M</i> (<i>SE</i>)	811 ± 10	808 ± 11	809 ± 11	816 ± 10
Neutral target	<i>M</i> (<i>SE</i>)	723 ± 9	747 ± 9	730 ± 9	735 ± 9

Table 3. Parameter estimates for the population-level effects of the maximal Bayesian model including random intercepts and slopes per participant. The beta-values of the parameters (*b*), estimated error (EE) and credible intervals (CI) are shown (Experiment 2).

Predictor	<i>b</i> (EE)	95% CI
Intercept	.76 (.01)	[.74, .78]
Incongruent target	.04 (.01)	[.02, .05]
Human distractor	.00 (.01)	[-.02, .01]
Object distractor	.00 (.01)	[-.02, .01]
Robot distractor	.00 (.01)	[-.02, .01]
Incongruent target x human distractor	-.01 (.01)	[-.01, .04]
Incongruent target x object distractor	.00 (.01)	[-.02, .02]
Incongruent target x robot distractor	.00 (.01)	[-.02, .02]

nificant main effect of target or distractors, nor any significant interaction effects. Overall, the participants' performance on the task was very accurate again (93.29%).

Reaction times. The repeated measures ANOVA on the reaction time data revealed a main effect of target, consistent with the expected Stroop interference in the incongruent condition of the task: $F(1, 50)=70.31, p<.001, \eta^2=.06$. Again, this showed that the task worked as expected. The target x distractor interaction was not significant: $F(3, 150)= 0.36, p=.78, \eta^2 =.0003$. Planned contrasts were computed using estimated marginal means. No contrast of interest reached significance: there was no difference between human faces and flower images in the incongruent condition: $t(300)= .094, p=.92$. The mean reaction times and standard errors are summarized in Table 2 and the Stroop interference scores are visualized in Figure 4.

Bayesian regression analysis (exploratory). Given the results of Experiment 2, we explored the extent to which our data provided compelling evidence for the null hypothesis (no enhanced Stroop effect when human faces are presented compared to the control flower condition) by using a Bayesian regression modelling approach {brms} package in R and Stan (Version 2.9.0, Bürkner, 2016), as the null cannot be confirmed with Frequentist statistics.

Following Balota and Yap (2011), we fitted an ex-gaussian distribution to data, as the response shows a strong right-skew (Figure 5). The ex-gaussian distribution is the convolution of the normal and exponential distributions and has been shown to provide a good fit to reaction time data (Balota & Yap, 2011). We included target word and distractor type as fixed effects predictors and included random intercepts and random slopes for each participant in a maximal random effects structure. The same weakly informative prior was applied to all variables, with a Student's

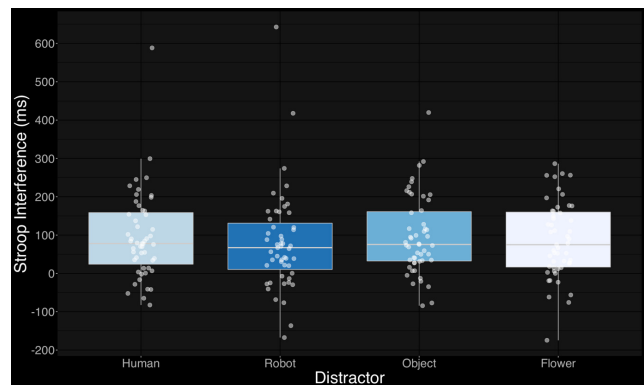


Figure 4. The mean Stroop interference scores (incongruent – neutral conditions) for each of the distractor categories in Experiment 2.

t-distribution of 3 degrees of freedom, a mean of 0 and a scale of 1. We used the default number of 4 Markov chains, each with 4000 iterations and a warm-up of 1000. This model converged, as supported by R-hat values below 1.01. We report the estimate (*b*), estimated error (EE) and the 95% credible interval in Table 3. The reaction time data was pre-processed in the same way as outlined in the data analysis section of Experiment 1.

To decide on the acceptance or rejection of a parameter null value we followed the approach outlined by Kruschke and colleagues (2018). Here, a range of plausible values are considered (indicated by the highest density interval (HDI) of the posterior distribution) and how they relate to a region of practical equivalence (ROPE) around null (Kruschke, 2018). The ROPE thus describes effects that are so small that they can be considered meaningless. In determining

the ROPE range, we set the limits following the procedure based on half of what we consider a small effect (Kruschke, 2018). A small effect in our first experiment was an average difference of 47ms between the incongruent social and incongruent control distractor, compared to a difference in 34ms in Conty and colleagues' task and 41ms in Chevallier and colleagues' version (2012, 2013). Choosing the most conservative small effect, we set the ROPE limits to [-.017, .017].

As depicted in Figure 6, the ROPE approach here does not offer a straightforward decision on the null hypothesis, even though zero is included in the range of credible parameter values, a small part of the HDI lies outside of the ROPE region for the effect of interest (slower reaction times for human distractors in the incongruent condition).

In summary, in defining our Bayesian regression model, we have increased the uncertainty of our estimates by including more random variance in the form of subject-level random effects. This increased uncertainty is expressed in Figure 5. Based on the ROPE analysis, we cannot definitively support the null hypothesis. However, considering that zero is contained in the 95% interval of credible values of the parameter's posterior distribution, the evidence for an effect is not very strong, and if real, goes in the opposite direction: -10ms [-10, 40].

General discussion

Across two experiments, we investigated how distracting faces with varying degrees of social salience were during a classic Stroop paradigm. Contrary to predictions derived from the fast track modulator model by Senju and Johnson (2009), and previous studies demonstrating robust attentional capture by task-irrelevant faces, we did not consistently observe the most salient social cues (human faces) leading to greater interference on the Stroop task. While we report a marginally significant interaction in Experiment 1, suggesting stronger distractibility of human faces in the incongruent condition, we caution interpretation of this finding, as we conducted our analysis on a smaller participant sample than planned. Thus, we reran our experiment with sufficient power, where we also used a larger number of unique distractor images. While we again observed the predicted general Stroop effect, the target by distractor interaction disappeared. Bayesian reanalysis of the data does not exclude the possibility of the human distractors influencing reaction times more than the neutral control distractors in the incongruent condition. However, this small predicted effect is likely not very strong. Overall, our findings contradict those reported by Conty and colleagues (2010) and Chevallier and colleagues (2013), who both found task-irrelevant social cues automatically captured attention. While their findings provided empirical evidence for the fast-track modulator model, which posits that social cues should exogenously and automatically engage attention, we don't see convincing evidence for this from our study. Our results not only appear counter-intuitive given the previous studies this work was based on, but also within the wider context of the literature documenting the reward value of social cues (Chevallier et al., 2012; Williams et al., 2019; Williams & Cross, 2018).

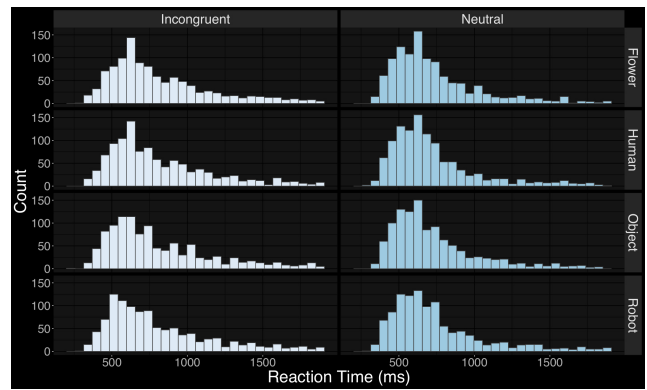


Figure 5. Distribution of the reaction times for each experimental condition (Experiment 2).




Figure 6. The region of practical equivalence (with zero) is shaded in gray. The effect of interest (the incongruent target with the human distractor image) is marked in dark blue as undecided (Experiment 2).

However, empirical evidence for social distractors always capturing attention is less convincing than the two studies by Conty and colleagues (2010) and Chevallier and colleagues (2013) suggest. A conceptual extension of their task from the lab of Hietanen, Myllyneva, Helminen and Lyyra (2016) failed to replicate the enhanced Stroop effect by direct gaze in a real-life version of the task. In their study, a confederate was looking at participants directly above a screen, which displayed a colour-matching version of the Stroop task. Hietanen and colleagues (2016) found a main effect of direct gaze speeding up the RTs of the participants, as compared to averted gaze. The authors reconcile their contradictory findings by relating them to the higher arousal produced by their stimuli: eye contact with a real person should be more engaging than pictorial representations thereof. In so-called low arousal contexts, they argue, salient social cues *should* recruit attentional resources and interfere with participants' performance on cognitive tasks. In our experiments, even in a context that Hietanen and colleagues (2016) describe as "low arousal", it is most probable that any social salience effect is practically equivalent to zero.

How can our results then be explained? Of course, the stimuli we presented were more complex than those used

in the original studies, so it is possible that the eye-contact effect only holds in (more) simplified contexts. The eye region in our stimulus set appeared smaller than in the original experiments, due to it taking up a smaller percentage of pixels in our distractor images. While the eye region itself was smaller, all our social stimuli (the human, robot and object faces) depicted direct gaze and a frontally oriented face. They only varied in their potential as a social interaction partner. So, if the eye-contact effect were to hold, we should have seen a consistent difference between our most salient social stimuli with direct eye gaze (the human faces) and the neutral control condition (flowers). The fact that our data did not support this pattern is especially surprising given that past studies examining direct gaze have also used full-face stimuli in similar, cognitively demanding tasks (Burton et al., 2009; Conty, Russo, et al., 2010).

A close look at the social attentional capture literature reveals a variety of methodological issues and contradicting findings across studies investigating faces and facial features as task-irrelevant distractors. Many studies report effects based on very small samples (some as small as 8 participants per experiment; Ariga & Arihara, 2017; Miyazaki et al., 2012; Sato & Kawahara, 2015, make bold statements based on modest statistical evidence (“the three-way interaction approached significance, $F(2,76) = 2.46, p < .10$ ”, p. 1103, Hietanen et al., 2016) or use small sets of distractor images which are repeated across many experimental trials (Bindemann et al., 2007; Theeuwes & Van der Stigchel, 2006). Indeed, some of these problematic confounds have been highlighted and tested by Pereira and colleagues (2019, 2020).

Pereira and colleagues (2019, 2020) systematically controlled for each known confound in the social attentional literature, including the perceived attractiveness of stimuli, low-level features and a list of other stimulus properties. In their studies, the authors utilized the dot-probe paradigm, with faces, houses and scrambled distractor images as task-irrelevant cues. The targets appeared with an equal likelihood at six different locations. Pereira and colleagues found across multiple experiments that faces did not reliably draw attention to their cued location, as indexed by participants' reaction time. In a follow-up Bayesian analysis on one of their experiments, the authors found evidence for the null hypothesis of no reaction time differences emerging for targets appearing at locations that were cued by faces or houses (Pereira et al., 2019). While a different task was used in these studies, the authors' findings closely align with ours: faces are not reliably capturing attention and impairing the performance on an unrelated cognitive task. Interestingly, in a direct replication of Bindemann and colleagues (2007), using less well-controlled stimuli, the authors were able to replicate the effect of attentional capture by task-irrelevant faces, providing convincing evidence for systematic confounds obscuring the true picture in the existing literature.

More evidence for the variable nature of findings on automatic attentional biasing by social cues comes from a series of experiments by Framorando and colleagues (2016), who, similar to Hietanen and colleagues (2016), also failed to replicate attentional capture by direct gaze, when faces were presented in a stare-in-crowd task paradigm. Based

on previous literature on this effect, one should expect that faces with direct gaze should be more distracting than faces with averted gaze. The authors found that straight gaze had a facilitatory effect when it was part of the target of the task, not a task-irrelevant distractor cue. These findings were later extended by the same authors, emphasizing again the task-dependent nature of directly gazing faces, which in this study hinged on the social or non-social nature of the task (2018). These empirical findings echo an fMRI study by Pessoa and colleagues (2002), who investigated attentional capture by *emotional* facial cues. Here, like the fast-track modulator model, a popular theory suggests that a subcortical route gives preference to the processing of emotional facial cues. However, the authors found that brain regions implicated in emotion perception were only active when participants were able to attend to the emotional facial cues, and these same brain regions were not differentially modulated when participants were engaged in a cognitively demanding task. This, the authors conclude, means that attentional resources are in fact necessary to allow the neural processing of emotional facial cues.

While we can reconcile our results with these studies, one may still wonder why social cues, which are thought to be inherently rewarding, failed to engage participants in our experiments in the expected manner (Anderson, 2016). Speaking to this, recent findings on reward-related distractors impairing participants' performance have also called this intuitive hypothesis into question (Rusz et al., 2019). A new meta-analysis suggests that the effect size of studies on reward-related distraction is small, and that findings across reviewed studies are highly variable, with reverse results not being uncommon (Rusz et al., 2020). This dovetails with the contradictory results we have found in the literature of social attentional biasing and which have also been addressed by Pereira and colleagues (2020).

Of course, based on this small number of empirical studies, we do not wish to claim that salient social cues, such as faces, never capture automatic attention in any context. Indeed, there is mounting evidence that overt attention (i.e. eye saccades towards social cues), as opposed to covert attention, which is measured by manual reaction time, is consistently directed towards the eye region of faces (DeAngelus & Pelz, 2009; Hayward et al., 2017; Pereira et al., 2020). Still, we do wish to challenge the putative fast track modulator model and speculate that when faces are presented as task-irrelevant distractors, they may not be salient enough to draw attention and cognitive processing resources away from the task at hand. Furthermore, we question the suitability of the task as a “proxy for social motivation”, as suggested by Chevallier and colleagues (p. 1649, 2013).

However, our findings should also be interpreted with the following limitations in mind: over the course of two experiments, we recruited from an ethnically diverse participant pool at the University of Glasgow, while presenting rather homogenous looking human faces, consisting exclusively of Caucasian individuals. Given that the studies we based our experiments on did not explicitly mention or measure this factor, we did not assess ethnic background in the short demographic survey preceding both studies. As such, we cannot test whether this aspect played a role in the missing en-

hanced Stroop interference effect for the human distractor images.

A further stimulus-based limitation was that in Experiment 1, distractors were not controlled by their mirror and presented twice. Thus, the repeat presentation could have led to a particularly memorable stimulus set. In Experiment 2, the unique distractors in the incongruent condition were controlled by their mirror images. Of course, on the other hand, the repeat presentation of distractor images is common practice in the social attentional capture literature (for example, a set of four unique human and pareidolic face images used for an experiment consisting of 450 trials, Ariga & Arihara, 2017). Takahashi and colleagues (2013) used stimuli with three unique identities over many trials, and only four unique stimuli in another study (Takahashi & Watanabe, 2015). Theeuwes et al. (2006) presented 12 unique distractor images across 96 trials. To put it differently, based on the conventions of the social attentional biasing literature, it is unlikely that we did not observe the expected effect due to the number of unique distractor images we presented.

Despite our best efforts to only include neutral faces, the emotional content of the social stimuli could not be controlled to a fine-grained degree, as it was limited by the design and availability of the robots and objects that were identified through our Google search. In the emotion rating experiment, which we undertook prior to Experiment 1, the robot faces were not rated as unambiguously neutral as the human faces, even after excluding the outliers. Human faces were selected from the neutral category of the Radboud and London faces database, so these stimuli would have contained inherently less variance in perceived emotionality than the robot and object faces. However, given the scarcity of frontally oriented and high-quality robot and object faces, we chose to operate within those constraints. Moreover, in comparison with other studies on social attentional biasing we were able to control for the following confounds (as outlined in Pereira et al., 2020): size and shape of the distractors, luminance and contrast, distance from fixation, the internal configuration of facial features of the human, robot and object images (i.e. a comparable set of features including eyes, a nose and a mouth in most of the images), as well as the task context.

While this set of experiments constitutes a conceptual extension to face stimuli, rather than a direct replication of the eye contact effect, we kept most other aspects of the experimental procedure identical to the studies we modelled our task on. Based on these studies and the facial attentional capture literature, we would have expected that human faces would be most salient, regardless of the small modifications we made. Indeed, keeping in mind recent calls for more generalisation efforts in psychological science (Yarkoni, 2016), we feel that a conceptual replication adds crucial insight to the field of motivated cognition. Further to the arguments we presented, our question and approach directly relate to the conceptualized fast-track modulator model: we tested and failed to support Chevallier and colleagues' (2013) hypothesis that this effect should generalize to other social cues – like faces – as well.

For future research, our findings have important implications: many researchers in human-robot interaction (HRI)

lament the absence of robust behavioural tasks to assess social interactions with robots, especially regarding changes in social motivation towards them (Baxter et al., 2016; Eyszel & Kuchenbrandt, 2012; Henschel & Cross, 2020). A few research groups have successfully adapted cognitive tasks for HRI, for example the inversion effect (to examine anthropomorphism), and the Posner gaze-cueing paradigm (Wykowska et al., 2014; Zlotowski & Bartneck, 2013). Yet, behavioural tasks that reliably assess social motivation towards robots are still scarce. Based on our findings, a suitable point of departure for future generations of social robotics researchers could be to examine overt attention in preferential looking paradigms or saccadic choice tasks, utilizing eye tracking technology (Crouzet & Thorpe, 2010; Fletcher-Watson et al., 2008), as these effects appear robust (Hayward et al., 2017). Another option could be to implement more natural social interaction tasks and measure attentional engagement and shifts in a similar manner as Hayward and colleagues in their conversational paradigm, in which participants' eye gaze behaviour was recorded with spyglasses and cameras (2017). Interestingly, the authors found that the social attention of participants in a natural context was unrelated to their behaviour in the classic Posner gaze cueing task. Their findings also speak to recent calls in the HRI literature to implement more natural, embodied experiments with robots to test changes in attitudes, behaviours and neural correlates in a more ecologically valid context (Henschel et al., 2020).

On a more fundamental level, one should reflect on the issue of small effect sizes to be expected in experimental psychology (Funder & Ozer, 2019; Ramsey, 2020; Schäfer & Schwarz, 2019). Based on the insights of recent large scale replication projects, we can be fairly certain that many established effects in the literature are much smaller than initially presented, if replicable at all (Camerer et al., 2018). One should then question what the smallest effect size is that one would consider interesting. Going forward, researchers should aim to conduct well-powered direct replications and consider expected effect sizes before adapting social motivation paradigms for HRI.

When Arcimboldo originally painted his whimsical portraits in the late 16th century, little did he know that machines today would be endowed with facial features to evoke illusory socialness – a simple, yet effective trick, corroborated by data that show that mechanical and screen-based robot faces are rated as humanlike, friendly, intelligent or in some cases, as uncanny (Chesher & Andreollo, 2020; Kalgina et al., 2018; Phillips et al., 2018; Vallverdú & Trovato, 2016). As our surroundings become increasingly populated by a variety of artificial agents (including robots and virtual agents), an important aim will be to probe how different types of faces are processed, and what we might learn about humans' intrinsic social motivation toward artificial agents' faces (Geiger & Balas, 2020).

Contributions

- Contributed to conception and design: AH, ESC
- Contributed to acquisition of data: AH, HB

- Contributed to analysis and interpretation of data: AH, HB, ESC
- Drafted and/or revised the article: AH, ESC, HB
- Approved the submitted version for publication: AH, ESC, HB

Acknowledgments

We would like to thank Lauren Colbert for her assistance with the data collection for Experiment 2 and Prof. Guillaume Thierry and Dr. Kami Koldewyn, who helped improve the design of these studies with their valuable feedback. We also want to express our gratitude towards Dr. Andrew Milne and Eline Smit for sharing their insights into Bayesian regression analysis and to Dr. Andrew Milne and Te-Yi Hsieh for their valuable comments on the final draft of the manuscript.

Funding information

This work has received funding from the European Research Council (ERC) under the European Union's Horizon2020 research and innovation programme (grant agree-

ment number ERC-2015-StG-677270-SOCIALROBOTS to ESC).

Supplementary Materials

Three text files which describe additional analyses are available.

Study 1: Emotion rating online validation study

Study 2: Agency and experience ratings

Study 3: Pre-processing of response times

Competing interests

The authors declare no competing interests.

Data accessibility statement

The pre-registration, the analysis scripts, the data and the stimulus material (available upon request) can be found via the Open Science Framework: <https://osf.io/xyz4m/>

Submitted: January 19, 2020 PST, Accepted: November 13, 2020 PST



This is an open-access article distributed under the terms of the Creative Commons Attribution 4.0 International License (CC-BY-4.0). View this license's legal deed at <http://creativecommons.org/licenses/by/4.0> and legal code at <http://creativecommons.org/licenses/by/4.0/legalcode> for more information.

REFERENCES

- Anderson, B. A. (2016). Social reward shapes attentional biases. *Cognitive Neuroscience*, 7(1–4), 30–36. <https://doi.org/10.1080/17588928.2015.1047823>
- Ariga, A., & Arihara, K. (2017). Visual attention is captured by task-irrelevant faces, but not by pareidolia faces. *2017 9th International Conference on Knowledge and Smart Technology: Crunching Information of Everything, KST 2017*, 1, 266–269. <https://doi.org/10.1109/KST.2017.7886105>
- Arslan, R. C., Walther, M. P., & Tata, C. S. (2019). formr: A study framework allowing for automated feedback generation and complex longitudinal experience-sampling studies using R. *Behavior Research Methods*, 1–37. <https://doi.org/10.3758/s13428-019-01236-y>
- Balota, D. A., & Yap, M. J. (2011). Moving beyond the mean in studies of mental chronometry: The power of response time distributional analyses. *Current Directions in Psychological Science*, 20(3), 160–166. <https://doi.org/10.1177/0963721411408885>
- Baxter, P., Kennedy, J., Senft, E., Lemaignan, S., & Belpaeme, T. (2016). From characterising three years of HRI to methodology and reporting recommendations. *ACM/IEEE International Conference on Human-Robot Interaction, 2016-April*, 391–398. <https://doi.org/10.1109/HRI.2016.7451777>
- Bindemann, M., Burton, A. M., Langton, S. R. H., Schweinberger, S. R., & Doherty, M. J. (2007). The control of attention to faces. *Journal of Vision*, 7(10), 1–8. <https://doi.org/10.1167/7.10.15>
- Bjornsdottir, R. T., & Rule, N. O. (2017). The Visibility of Social Class From Facial Cues. *Journal of Personality and Social Psychology*, 113(4), 530–546. <https://doi.org/10.1037/pspa0000091.supp>
- Bradley, M. M., & Lang, P. P. J. (1999). Affective Norms for English Words (ANEW): Instruction Manual and Affective Ratings. *Psychology, Technical(C-1)*, 0. <https://doi.org/10.1109/MIC.2008.114>
- Bubic, A., Susac, A., & Palmovic, M. (2014). Keeping our eyes on the eyes: The case of Arcimboldo. *Perception*, 43(5), 465–468. <https://doi.org/10.1068/p7671>
- Bürkner, P. C. (2016). Package ‘brms.’
- Burra, N., Framorando, D., & Pegna, A. J. (2018). Early and late cortical responses to directly gazing faces are task dependent. *Cognitive, Affective and Behavioral Neuroscience*, 18(4), 796–809. <https://doi.org/10.3758/s13415-018-0605-5>
- Burton, A. M., Bindemann, M., Langton, S. R. H., Schweinberger, S. R., & Jenkins, R. (2009). Gaze Perception Requires Focused Attention: Evidence From an Interference Task. *Journal of Experimental Psychology: Human Perception and Performance*, 35(1), 108–118. <https://doi.org/10.1037/0096-1523.35.1.108>
- Camerer, C. F., Dreber, A., Holzmeister, F., Ho, T. H., Huber, J., Johannesson, M., Kirchler, M., Nave, G., Nosek, B. A., Pfeiffer, T., Altmeld, A., Buttrick, N., Chan, T., Chen, Y., Forsell, E., Gampa, A., Heikensten, E., Hummer, L., Imai, T., ... Wu, H. (2018). Evaluating the replicability of social science experiments in Nature and Science between 2010 and 2015. *Nature Human Behaviour*, 2(9), 637–644. <https://doi.org/10.1038/s41562-018-0399-z>
- Chesher, C., & Andreollo, F. (2020). Robotic Faciality: The Philosophy, Science and Art of Robot Faces. *International Journal of Social Robotics, January*. <https://doi.org/10.1007/s12369-020-00623-2>
- Chevallier, C., Hugué, P., Happé, F., George, N., & Conty, L. (2013). Salient social cues are prioritized in autism spectrum disorders despite overall decrease in social attention. *Journal of Autism and Developmental Disorders*, 43(7), 1642–1651. <https://doi.org/10.1007/s10803-012-1710-x>
- Chevallier, C., Kohls, G., Troiani, V., Brodtkin, E. S., & Schultz, R. T. (2012). The social motivation theory of autism. *Trends in Cognitive Sciences*, 16(4), 231–238. <https://doi.org/10.1016/j.tics.2012.02.007>
- Conty, L., Gimmig, D., Belletier, C., George, N., & Hugué, P. (2010). The cost of being watched: Stroop interference increases under concomitant eye contact. *Cognition*, 115(1), 133–139. <https://doi.org/10.1016/j.cognition.2009.12.005>
- Conty, L., Russo, M., Loehr, V., Hugueville, L., Barbu, S., Hugué, P., Tijus, C., & George, N. (2010). The mere perception of eye contact increases arousal during a word-spelling task. *Social Neuroscience*, 5(2), 171–186. <https://doi.org/10.1080/17470910903227507>
- Crouzet, S. M., & Thorpe, S. J. (2010). Fast saccades toward faces: Face detection in just 100 ms. *Journal of Vision*, 10(2010), 1–17. <https://doi.org/10.1167/10.4.16.Introduction>

- DeAngelus, M., & Pelz, J. B. (2009). Top-down control of eye movements: Yarbus revisited. *Visual Cognition*, 17(6–7), 790–811. <https://doi.org/10.1080/135062809.02793843>
- DeBruine, L., & Jones, B. (2017). *Face Research Lab London set*. https://figshare.com/articles/dataset/Face_Research_Lab_London_Set/5047666
- DiSalvo, C. F., & Gemperle, F. (2003). From Seduction to Fulfillment: The Use of Anthropomorphic Form in Design. *Proceedings of the International Conference on Designing Pleasurable Products and Interfaces*, 67–72.
- DiSalvo, C. F., Gemperle, F., Forlizzi, J., & Kiesler, S. (2002). All robots are not created equal: The design and perception of humanoid robot heads. *Proceedings of the Conference on Designing Interactive Systems: Processes, Practices, Methods, and Techniques, DIS*, 321–326. <https://doi.org/10.1145/778712.778756>
- Eyssel, F., & Kuchenbrandt, D. (2012). Social categorization of social robots: Anthropomorphism as a function of robot group membership. *British Journal of Social Psychology*, 51(4), 724–731. <http://doi.org/10.1111/j.2044-8309.2011.02082.x>
- Fletcher-Watson, S., Findlay, J. M., Leekam, S. R., & Benson, V. (2008). Rapid detection of person information in a naturalistic scene. *Perception*, 37(4), 571–583. <https://doi.org/10.1068/p5705>
- Framorando, D., George, N., Kerzel, D., & Burra, N. (2016). Straight gaze facilitates face processing but does not cause involuntary attentional capture. *Visual Cognition*, 24(7–8), 381–391. <https://doi.org/10.1080/13506285.2017.1285840>
- Funder, D. C., & Ozer, D. J. (2019). Evaluating Effect Size in Psychological Research: Sense and Nonsense. *Advances in Methods and Practices in Psychological Science*, 2(2), 156–168. <https://doi.org/10.1177/2515245919847202>
- Geiger, A., & Balas, B. (2020). *Not quite human, not quite machine: Electrophysiological responses to robot faces*.
- Gray, H. M., Gray, K., & Wegner, D. M. (2007). Dimensions of mind perception. *Science*, 315(5812), 619–619. <https://doi.org/10.1126/science.1134475>
- Guido, G., Pichierri, M., Pino, G., & Natarajan, R. (2019). Effects of face images and face pareidolia on consumers' responses to print advertising: An empirical investigation. *Journal of Advertising Research*, 59(2), 219–231. <https://doi.org/10.2501/JAR-2018-030>
- Hayward, D. A., Voorhies, W., Morris, J. L., Capozzi, F., & Ristic, J. (2017). Staring Reality in the Face: A Comparison of Social Attention Across Laboratory and Real World Measures Suggests Little Common Ground. *Canadian Journal of Experimental Psychology*, 71(3), 212–225. <https://doi.org/10.1037/cp0000117>
- Henschel, A., & Cross, E. S. (2020). No evidence for enhanced likeability and social motivation towards robots after synchrony experience. *Interaction Studies*, 21(1), 7–23. <https://doi.org/10.1075/is.19004.hen>
- Henschel, A., Hortensius, R., & Cross, E. S. (2020). Social Cognition in the Age of Human – Robot Interaction. *Trends in Neurosciences*, 43(6), 1–12. <https://doi.org/10.1016/j.tins.2020.03.013>
- Hessels, R. S. (2020). How does gaze to faces support face-to-face interaction? A review and perspective. *Psychonomic Bulletin and Review*. <https://doi.org/10.3758/s13423-020-01715-w>
- Hietanen, J. K., Myllyneva, A., Helminen, T., & Lyyra, P. (2016). The Effects of Genuine Eye Contact on Visuospatial and Selective Attention. *Journal of Experimental Psychology*, 41(6), 573–575. <https://doi.org/10.1007/s001080050562>
- Isaac, L., Vrijssen, J. N., Eling, P., Van Oostrom, I., Speckens, A., & Becker, E. S. (2012). Verbal and facial-emotional stroop tasks reveal specific attentional interferences in sad mood. *Brain and Behavior*, 2(1), 74–83. <https://doi.org/10.1002/brb3.38>
- Kalegina, A., Schroeder, G., Allchin, A., Berlin, K., & Cakmak, M. (2018). Characterizing the Design Space of Rendered Robot Faces. *ACM/IEEE International Conference on Human-Robot Interaction*, 96–104. <https://doi.org/10.1145/3171221.3171286>
- Kruschke, J. K. (2018). Rejecting or accepting parameter values in Bayesian estimation. *Advances in Methods and Practices in Psychological Science*, 1(2), 270–280. <https://doi.org/10.1177/2515245918771304>
- Langner, O., Dotsch, R., Bijlstra, G., Wigboldus, D. H., Hawk, S. T., & Van Knippenberg, A. D. (2010). Presentation and validation of the Radboud Faces Database. *Cognition and Emotion*, 24(8), 1377–1388.
- Langton, S. R. H., Law, A. S., Burton, A. M., & Schweinberger, S. R. (2008). Attention capture by faces. *Cognition*, 107(1), 330–342. <https://doi.org/10.1016/j.cognition.2007.07.012>
- Lawrence, M. A. (2016). *ez: Easy analysis and visualization of factorial experiments*. R package version, 4(2).

- Lenth, R., Singmann, H., Love, J., Buerkner, P., & Herve, M. (2019). *emmeans: Estimated Marginal Means, aka Least-Squares Means*. (Version 1.3. 4).
- Liu, J., Li, J., Feng, L., Li, L., Tian, J., & Lee, K. (2014). Seeing Jesus in toast: Neural and behavioral correlates of face pareidolia. *Cortex*, *53*(1), 60–77. <https://doi.org/10.1016/j.cortex.2014.01.013>
- 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. <https://doi.org/10.1177/0956797610388044>
- MacLeod, C. M., & MacDonald, P. A. (2000). Interdimensional interference in the Stroop effect: Uncovering the cognitive and neural anatomy of attention. *Trends in Cognitive Sciences*, *4*(10), 383–391. [https://doi.org/10.1016/S1364-6613\(00\)01530-8](https://doi.org/10.1016/S1364-6613(00)01530-8)
- Martinez-Conde, S., Conley, D., Hine, H., Kropf, J., Tush, P., Ayala, A., & Macknik, S. L. (2015). Marvels of illusion: Illusion and perception in the art of Salvador Dali. *Frontiers in Human Neuroscience*, *9*(SEPTEMBER), 1–12. <https://doi.org/10.3389/fnhum.2015.00496>
- Miyazaki, Y., Wake, H., Ichihara, S., & Wake, T. (2012). Attentional Bias to Direct Gaze in a Dot-Probe Paradigm. *Perceptual and Motor Skills*, *114*(3), 1007–1022. <https://doi.org/10.2466/21.07.24.pms.114.3.1007-1022>
- Omer, Y., Sapir, R., Hatuka, Y., & Yovel, G. (2019). What Is a Face? Critical Features for Face Detection. *Perception*, *48*(5), 437–446. <https://doi.org/10.1177/0301006619838734>
- Palmer, C. J., & Clifford, C. W. G. (2020). Face Pareidolia Recruits Mechanisms for Detecting Human Social Attention. *Psychological Science*, *095679762092481*. <https://doi.org/10.1177/0956797620924814>
- Pavlova, M. A., Galli, J., Pagani, F., Micheletti, S., Guerreschi, M., Sokolov, A. N., Fallgatter, A. J., & Fazzi, E. M. (2018). Social cognition in down syndrome: Face tuning in face-like non-face images. *Frontiers in Psychology*, *9*(DEC), 1–9. <https://doi.org/10.3389/fpsyg.2018.02583>
- Peirce, J. W. (2007). PsychoPy-psychophysics software in Python. *Journal of Neuroscience Methods*, *162*(1–2), 8–13. <https://doi.org/10.1016/j.jneumeth.2006.11.017>
- Pereira, E., Birmingham, E., & Ristic, J. (2019). Contextually-Based Social Attention Diverges across Covert and Overt Measures. *Vision*, *26*(1), 46–46. <https://doi.org/10.3390/vision3020029>
- Pereira, E., Birmingham, E., & Ristic, J. (2020). The eyes do not have it after all? Attention is not automatically biased towards faces and eyes. *Psychological Research*, *84*(5), 1407–1423. <https://doi.org/10.1007/s00426-018-1130-4>
- Pessoa, L., McKenna, M., Gutierrez, E., & Ungerleider, L. G. (2002). Neural processing of emotional faces requires attention. *Proceedings of the National Academy of Sciences of the United States of America*, *99*(17), 11458–11463. <https://doi.org/10.1073/pnas.172403899>
- Phillips, E., Zhao, X., Ullman, D., & Malle, B. F. (2018). *What is Human-like?: Decomposing Robots' Human-like Appearance Using the Anthropomorphic roBOT (ABOT) Database*. *18*, 105–113.
- Ramsey, R. (2020). A call for greater modesty in psychology and cognitive neuroscience. *PsyArXiv*, 1–23. <https://doi.org/10.31234/osf.io/hf5sv>
- Revelle, W. (2018). *psych: Procedures for personality and psychological research*. Northwestern University.
- Ro, T., Russell, C., & Lavie, N. (2001). Changing faces: A detection advantage in the flicker paradigm. *Psychological Science*, *12*(1), 94–99. <https://doi.org/10.1111/1467-9280.00317>
- Robertson, D. J., Jenkins, R., & Burton, A. M. (2017). Face detection dissociates from face identification. *Visual Cognition*, *25*(7–8), 740–748. <https://doi.org/10.1080/13506285.2017.1327465>
- Rusz, D., Bijleveld, E., & Kompier, M. (2019). Do Reward-Related Distractors Impair Cognitive Performance? Perhaps Not. *Collabra: Psychology*, *5*(1), 10. <https://doi.org/10.1525/collabra.169>
- Rusz, D., Le Pelley, M., Kompier, M. A. J., Mait, L., & Bijleveld, E. (2020). Reward-driven distraction: A meta-analysis. *Journal of Chemical Information and Modeling*. <https://doi.org/10.1017/CBO9781107415324.004>
- Sato, S., & Kawahara, J. I. (2015). Attentional capture by completely task-irrelevant faces. *Psychological Research*, *79*(4), 523–533. <https://doi.org/10.1007/s00426-014-0599-8>

- Schäfer, T., & Schwarz, M. A. (2019). The meaningfulness of effect sizes in psychological research: Differences between sub-disciplines and the impact of potential biases. *Frontiers in Psychology, 10*(APR), 1–13. <https://doi.org/10.3389/fpsyg.2019.00813>
- Schilbach, L., Eickhoff, S. B., Cieslik, E., Shah, N. J., Fink, G. R., & Vogeley, K. (2011). Eyes on me: An fMRI study of the effects of social gaze on action control. *Social Cognitive and Affective Neuroscience, 6*(4), 393–403. <https://doi.org/10.1093/scan/nsq067>
- Senju, A., & Johnson, M. H. (2009). The eye contact effect: mechanisms and development. *Trends in Cognitive Sciences, 13*(3), 127–134. <https://doi.org/10.1016/j.tics.2008.11.009>
- Simmons, J. P., Nelson, L. D., & Simonsohn, U. (2012). A 21 Word Solution. *SSRN Electronic Journal, 1*–4. <https://doi.org/10.2139/ssrn.2160588>
- Singmann, H., Bolker, B., Westfall, J., Aust, F., & Ben-Shachar, M. S. (2019). *Afex: Analysis of factorial experiments*. R package version 0.25-1.
- Stroop R. (1935). Studies of Interference in Serial Verbal Reactions. *Journal of Experimental Psychology, 18*(6), 643–661.
- Takahashi, K., & Watanabe, K. (2013). Gaze cueing by pareidolia faces. *I-Perception, 4*(8), 490–492. <http://doi.org/10.1068/i0617sas>
- Takahashi, K., & Watanabe, K. (2015). Seeing objects as faces enhances object detection. *I-Perception, 6*(5), 1–14. <https://doi.org/10.1177/2041669515606007>
- Theeuwes, J., & Van der Stigchel, S. (2006). Faces capture attention: Evidence from inhibition of return. *Visual Cognition, 13*(6), 657–665. <https://doi.org/10.1080/13506280500410949>
- Vallverdú, J., & Trovato, G. (2016). Emotional affordances for human–robot interaction. *Adaptive Behavior, 24*(5), 320–334. <https://doi.org/10.1177/1059712316668238>
- van Honk, J., Tuiten, A., de Haan, E., van de Hout, M., & Stam, H. (2001). Attentional biases for angry faces: Relationships to trait anger and anxiety. *Cognition and Emotion, 15*(3), 279–297. <https://doi.org/10.1080/02699930126112>
- Van Honk, J., Tuiten, A., Van Den Hout, M., Koppeschaar, H., Thijssen, J., De Haan, E., & Verbaten, R. (2000). Conscious and preconscious selective attention to social threat: Different neuroendocrine response patterns. *Psychoneuroendocrinology, 25*(6), 577–591. [https://doi.org/10.1016/S0306-4530\(00\)00011-1](https://doi.org/10.1016/S0306-4530(00)00011-1)
- Vuilleumier, P. (2002). Facial expression and selective attention. *Current Opinion in Psychiatry, 15*(3), 291–300. <https://doi.org/10.1097/00001504-200205000-00011>
- Willenbockel, V., Sadr, J., Fiset, D., Horne, G. O., Gosselin, F., & Tanaka, J. W. (2010). Controlling low-level image properties: The SHINE toolbox. *Behavior Research Methods, 42*(3), 671–684. <https://doi.org/10.3758/BRM.42.3.671>
- Williams, E. H., Cristino, F., & Cross, E. S. (2019). Human body motion captures visual attention and elicits pupillary dilation. *Cognition, 193*(January), 104029. <https://doi.org/10.1016/j.cognition.2019.104029>
- Williams, E. H., & Cross, E. S. (2018). Decreased reward value of biological motion among individuals with autistic traits. *Cognition, 171*(March 2017), 1–9. <https://doi.org/10.1016/j.cognition.2017.10.017>
- Willis, J., & Todorov, A. (2006). First impressions: Making up your mind after a 100-ms exposure to a face. *Psychological Science, 17*(7), 592–598. <https://doi.org/10.1111/j.1467-9280.2006.01750.x>
- Wodehouse, A., Brisco, R., Broussard, E., & Duffy, A. (2018). Pareidolia: Characterising facial anthropomorphism and its implications for product design. *Journal of Design Research, 16*(2), 83–98. <https://doi.org/10.1504/JDR.2018.092792>
- 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). <https://doi.org/10.1371/journal.pone.0094339>
- Yarbus, A. L. (1967). *Eye Movements and Vision* (B. Haigh, Trans.). Plenum Press.
- Yarkoni, T. (2016). The generalizability crisis. *PsyArXiv, 1*–26.
- Zhou, L. F., & Meng, M. (2019). Do you see the “face”? Individual differences in face pareidolia. *Journal of Pacific Rim Psychology, 13*(1). <https://doi.org/10.1017/prp.2019.27>
- Zlotowski, J., & Bartneck, C. (2013). The inversion effect in HRI: Are robots perceived more like humans or objects? *ACM/IEEE International Conference on Human-Robot Interaction, 365–372*. <https://doi.org/10.1109/HRI.2013.6483611>

SUPPLEMENTARY MATERIALS

Peer Review History

Download: https://collabra.scholasticahq.com/article/18521-faces-do-not-attract-more-attention-than-non-social-distractors-in-the-stroop-task/attachment/48491.docx?auth_token=MYiXBZW6q1ZFFEbCh1r

Study 1: Emotion rating online validation study

Download: https://collabra.scholasticahq.com/article/18521-faces-do-not-attract-more-attention-than-non-social-distractors-in-the-stroop-task/attachment/48494.docx?auth_token=MYiXBZW6q1ZFFEbCh1r

Study 2: Agency and experience ratings

Download: https://collabra.scholasticahq.com/article/18521-faces-do-not-attract-more-attention-than-non-social-distractors-in-the-stroop-task/attachment/49493.docx?auth_token=MYiXBZW6q1ZFFEbCh1r

Study 3: Pre-processing of response times

Download: https://collabra.scholasticahq.com/article/18521-faces-do-not-attract-more-attention-than-non-social-distractors-in-the-stroop-task/attachment/49494.docx?auth_token=MYiXBZW6q1ZFFEbCh1r
