

Differences in Working-Memory Capacity Modulate Top-down Control of Social Attention

Ali Momen
George Mason University
Fairfax, Virginia

Eva Wiese
George Mason University
Fairfax, Virginia

ABSTRACT

Gaze following, or our ability to attend to where others are looking can be top-down controlled by context information about the social relevance of the gaze signal. In particular, it has been shown that gaze signals are followed more strongly when the gazer is believed to have a mind with the ability to show intentional behavior (i.e., human) compared to being pre-programmed (i.e., robot). Perceiving human traits in nonhuman agents (i.e., anthropomorphism) occurs naturally in human-robot interaction, where it has positive effects on performance. It can also attenuate performance, if the robot is designed in a way that makes it hard to categorize as human or nonhuman (e.g., humanoid appearance), and inflicts additional working memory load due to categorical ambiguity. Here, we examine if gaze signals of ambiguous humanoid agents are followed differently than those of unambiguous human or robot agents, and to what extent gaze following is affected by individual differences in working memory capacity (WMC). We assume participants with high versus low WMC to be more capable of suppressing reflexive gaze following behaviors to the cued location in a counterpredictive paradigm (where targets appear with high likelihoods at uncued locations), particularly when being cued by humanoid gazers (top-down control, which requires cognitive resources). While the analysis showed no effect of categorical ambiguity on top-down control abilities overall, it revealed that participants with low WMC had weaker top-down control than participants with high WMC for the most ambiguous humanoid agent. The results are discussed with regard to the design of social agents and human-robot interaction.

CCS CONCEPTS

• **Applied computing** → *Psychology*;

KEYWORDS

Human-Robot Interaction, Gaze Cueing, Working Memory

ACM Reference Format:

Ali Momen and Eva Wiese. 2018. Differences in Working-Memory Capacity Modulate Top-down Control of Social Attention. In *TechMindSociety '18: APAScience '18: Technology, Mind, and Society, April 5–7, 2018, Washington, DC, USA*. ACM, New York, NY, USA, 6 pages. <https://doi.org/10.1145/3183654.3183686>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

TechMindSociety '18, April 5–7, 2018, Washington, DC, USA

© 2018 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-5420-2/18/04...\$15.00

<https://doi.org/10.1145/3183654.3183686>

1 INTRODUCTION

Human-robot interaction is a burgeoning field in which one of the main strategies for improving interactions between human and robot partners involves designing robot appearance and behavior in a way that makes the robot appear like a social companion [30, 31]. A special focus has been on equipping social robots with the ability to use non-verbal signals like facial expression or gaze direction to express their internal states and provide social reinforcement [5, 20]. Gaze direction can also be used to coordinate turn-taking in conversations [5], establish joint attention during shared tasks [27], and indicate what others are going to do next [13]. It is also a potent cue to warn others about potential dangers in the environment [7], and helps reducing ambiguity in conversations [2].

When examining the effect of observing others' gaze direction experimentally, a face is presented centrally on a computer screen that first gazes straight ahead, and then changes gaze direction to the left or right side of the screen. This gaze cue is then followed by a target that either appears at the gazed-at location (i.e., valid trial) or a gazed-away-from location (i.e., invalid trial) [25]. It has been shown that observing changes in others' gaze direction shifts the observer's attention to the gazed-at location, with the consequence that reactions to targets presented on valid trials are faster than on invalid trials. Differences in reaction times between invalid and valid trials represent the gaze-cueing effect [3, 12], and indicate how strongly the gaze signal is followed. Gaze cueing is strongest, when gaze direction is highly predictive of the target location (i.e., target appears at cued location in 80% of trials), but it can also be observed when gaze direction is non-predictive (i.e., target appears at cued location in 50% of the trials), or even counterpredictive of target location (i.e., target appears at uncued location in 80% of trials; see [12]). These findings suggest that it is difficult for participants to suppress gaze following (even when following the gaze cue is detrimental to task performance as with counterpredictive cues), which has been taken as strong evidence for the fact that gaze signals are followed reflexively, in a bottom-up manner [10]. Indeed, a wealth of studies showed that following gaze is so reflexive that it can be triggered by any stimulus that has eye-like configurations (e.g., eyes on a glove or apple, [26]), that it can be observed in infants as young as three-month old of age [16], and that it is undisturbed even under high workload [15]. At the same time, early studies failed to provide evidence that gaze following can be decreased, suppressed or enhanced (top-down controlled), based on factors like knowledge about the gazer (e.g., familiarity, trustworthiness, similarity) or interfering task goals (e.g., find target letter); see [3, 4, 12, 26].

However, recent studies provide evidence that the automaticity and degree to which gaze signals are followed can be top-down controlled, but only under specific circumstances [1, 9, 29, 32–34].

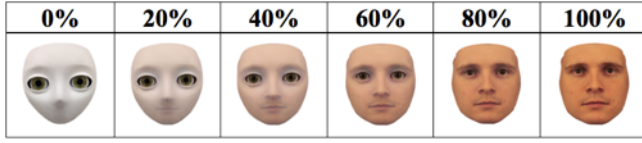


Figure 1: Morph spectrum ranging from 0% physical human-ness (on the left) to 100% physical humanness (on the right).

One way to examine how bottom-up and top-down processes interact during gaze following is by using a counterpredictive paradigm, where targets more likely appear at uncued locations (i.e., 80% of trials) than at the cued location (i.e., 20% of trials), and the task is to react as fast and accurately as possible to the presence, location or identity of the target [11]. If gaze following is a pure bottom-up process, knowledge about the counterpredictivity of the cue will not be used to orient away from the cued (but unlikely) location to the uncued (but likely) target location (i.e., reaction times would be shorter on valid versus invalid trials, and the difference score between invalid and valid trials would be positive). If gaze following can be top-down controlled, participants will use their knowledge about the counterpredictivity of the gaze cue and suppress reflexive orienting to the cued (but unlikely) location, and instead shift their attention to the uncued (but likely) target location (i.e., reaction times would be longer on valid versus invalid trials, and the difference score between invalid and valid trials would be negative). As a consequence, in counterpredictive cueing paradigms, the difference in reaction times between invalid and valid conditions can be used as a measure for the extent to which following gaze signals is top-down controlled: the more positive (negative) the difference in reaction times, the more pronounced the bottom-up (top-down) component is.

Using counterpredictive cueing, Friesen et al. [11] showed that gaze cues trigger reflexive (bottom-up) shifts of attention to the cued location and volitional (top-down modulated) shifts of attention to the predicted location, while orienting to arrow cues is always volitional. This finding suggests that although gaze following can be top-down controlled, it seems to be more reflexive than orienting to symbolic cues, probably due to its relevance for human-human interactions and human development [28]. Martini et al. [20] examined this hypothesis further by measuring if attentional orienting to human gaze cues is more reflexive (less top-down modulated) than gaze following triggered by non-human agents. The authors used a counterpredictive gaze cueing paradigm, in which participants were cued by morphed agents of varying physical human-likeness (spectrum from 100% physical robot to 100% physical human in 20% increments; see Figure 1). The authors hypothesized that the bottom-up orienting component would be most pronounced for the 100% human gazer (i.e., positive difference scores between invalid and valid trials), and then decrease as the physical humanness of the gazer decreases, accompanied by an increase in top-down control (i.e., difference scores between invalid and valid trials decrease along the spectrum and eventually turn negative). The results did not show the expected linear relationship between the gazer's humanness and the extent of bottom-up attentional orienting, but show an inverted U-shaped pattern instead:

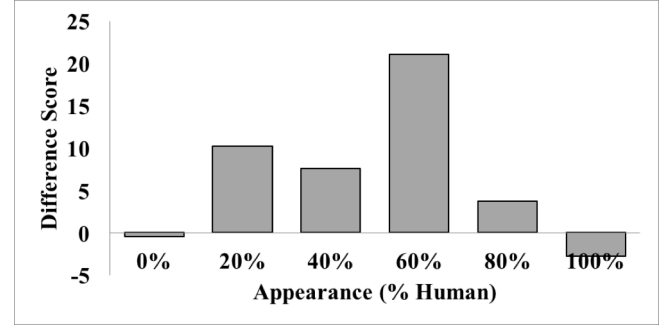


Figure 2: Extent of bottom-up and top-down attentional orienting along the human-robot morphing spectrum, as reported in Martini et al. [20]. Positive values indicate bottom-up orienting to the cued, but unlikely location, while negative values represent top-down controlled shifts of attention to the uncued, but likely location.

participants successfully managed to suppress reflexive shifts of attention to the cued (but unlikely) location in response to gaze cues sent by the 100% human and 100% robot agents (i.e., negative difference scores between invalid and valid trials), but were unable to do so for the other agents (i.e., positive difference scores between invalid and valid trials), with the strongest bottom-up effect for the 60% human morph; see Figure 2. Since exerting top-down control on attentional orienting relies on cognitive resources [23], the results suggest that processing agents with mixed human and robot characteristics drains cognitive resources, which are then not available for the top-down modulated suppression of reflexive gaze following to unlikely target locations.

The observation that agents of around 70% physical humanness seem to drain cognitive resources more strongly than other agents (i.e., most positive difference scores) can be explained within the framework of the uncanny valley theory, which states that agents that are human-like, but not perfectly human (i.e., 60-80% physical humanness) elicit feelings of eeriness [20, 24], potentially due to their categorical ambiguity (i.e., hard to categorize as human versus robot; [18, 21, 22]). In line with this observation, Mandell et al. [19] have shown that agents of about 70% humanness indeed drain cognitive resources on a sustained attention task, while none of the other agents on the human-robot morph had a comparable detrimental effect on cognitive resources.

1.1 Aim of Study

If the extent to which gaze following can be top-down controlled in a counterpredictive paradigm indeed depends on the availability of cognitive resources, individual differences in working memory capacity (WMC) should have a modulatory effect on bottom-up attentional orienting to the gaze signals, in particular to those of ambiguous agents. We hypothesize that categorically ambiguous agents will drain cognitive resources across participants, but that participants with high versus low WMC should still have enough cognitive resources to exert top-down control on gaze following and orient away from the cue to the predicted location. We also expect this effect to be strongest for agents whose humanness is

most ambiguous (i.e., morphs that fall into the uncanny valley at around 80% physical humanness). If that were true, participants with a high WMC should have less pronounced positive difference scores than participants with a low WMC (indicating that they tried to orient away from the cued to the predicted location), in particular for the 80% morph.

2 METHODS AND MATERIALS

2.1 Participants

37 undergraduate students (age: $M = 22.56$, $SD = 6.40$; 24 females) participated in the experiment for course credit. They reported normal or corrected-to-normal vision and gave informed consent at the beginning of the experiment. The Institutional Review Board (IRB) approved the stimuli and the experimental procedure. The experiment took about 40 minutes and participants were debriefed at the end of it.

2.2 Stimuli

The morphed stimuli were created with FantaMorph, which is a software tool that allows two agent images to be morphed into each other in order to create a spectrum of stimuli at specified increments. The two images morphed into each other for the current study was one that depicted the robot Meka (Meka Robotics) and another one that depicted the face of a male person (from the Karolinska Directed Emotional Faces database). The generated spectrum consisted of six agents of varying degrees of physical humanness: 0%, 20%, 40%, 60%, 80%, and 100%; see Figure 1. The images were the same images as the ones used by Martini et al. [20]. The images were 7.8"x 8.6" in size, were presented in color on a white background, and were aligned centrally on the horizontal axis of the screen. Images with averted gaze were created with Photoshop by shifting the iris and pupils of the original images to the left or right by 0.4°. The target stimuli to which participants responded (F and T) subtended 0.5"x 0.9", and were presented on the horizontal axis of the screen (14.7° away from center).

2.3 Procedure

Participants first received instructions and signed the informed consent. After that, they engaged in a block of 20 practice trials to familiarize themselves with the gaze-following procedure. Upon completing the practice block, the experiment began, and participants completed six blocks (one for each of the morphs) of 40 trials each (32 invalid, 8 valid), which resulted in a predictivity of 20% (i.e. counterpredictive). The order in which the blocks were presented was randomized to avoid potential biasing effects due to exposure to specific agents from the morph spectrum at the beginning of the experiment. To control for how strongly participants anthropomorphized the agents, participants first rated an agent at the beginning of a block in terms of its ability of having a mind ("Do you think the agent has a mind?" on a 8-point Likert scale: 1 = definitely not; 8 = definitely yes), and then performed the gaze following task with the same agent. The question was chosen based on previous research on anthropomorphism [14].

Each trial of gaze following began with a fixation cross that was presented for a random time interval of 700-1000 ms at the center of the screen, which was followed by the appearance of one of the

morphed agents looking straight ahead (i.e., mutual gaze) for 700-1000 ms. Afterwards, the agent shifted its gaze direction either to the left or right side of the screen (i.e., averted gaze; the illusion of an eye movement was created by presenting the agent image with averted gaze after the image with straight gaze). After a stimulus onset asynchrony (SOA) of 500 ms, one of the target letters (F or T) appeared on the screen (either at the valid or invalid location), and participants were asked to respond as quickly and accurately as possible to the identity of the target (i.e., discrimination task; see [10]) by pressing either the D or K key on a standard computer keyboard. The keys were marked with F and T stickers and the assignment of the keys was counterbalanced across participants (i.e., F/T was assigned to D/K in 50% of the trials and to K/D in the other 50%). Target letter and agent stimulus remained on screen until participants gave a response, or until a time-out criterion of 1200 ms was reached, whichever came first. After an intertrial interval (ITI) of 680 ms, the next trial started; see Figure 3.

After the gaze-following task was completed (total of 240 trials), individual WMC was measured using the Operation Span (Ospan), Symmetry Span (Sspan) and Reading Span (Rspan); see [8, 23]. In the Ospan, participants alternated between being presented with a mathematical operation they had to judge the validity of, and being presented with a to-be recalled letter. In the Sspan, participants alternated between being presented with the pattern of white and black squares in an 8x8 matrix in which they judged whether the pattern was symmetrical along the vertical axis, and being presented with a to-be-remembered position of a red dot on a 4x4 matrix. In the Rspan, participants had to judge the validity of a sentence while being presented with a to-be recalled letter. The scores for the Ospan and Rspan ranged from 0-30, and for the Sspan from 0-24. Total testing time was approximately 45 minutes.

2.4 Analysis

Three different dependent variables were collected during the experiment: mind ratings as a measure for ascribed humanness, difference scores between invalid and valid trials as a measure for the extent of exerted top-down control, and Ospan, Sspan and Rspan scores as a measure for individual WMC.

Mind ratings were analyzed via a univariate ANOVA with % physical humanness as a within-factor, and we expected to find significantly larger mind ratings for agents with higher levels of physical humanness. Behavioral data obtained from the gaze-following task was analyzed by first averaging the reaction times for valid and invalid trials (per agent), and then calculating the difference in reaction times between the invalid and valid conditions (per agent). The difference scores were used as a measure for the extent of exerted top-down control on gaze following (i.e., negative scores indicate top-down control, positive scores indicate bottom-up orienting), and were analyzed with a univariate ANOVA with % physical humanness as within-participants factor. We expected to find a significant main effect of % physical humanness with less positive difference scores for the 100% human and robot agent than for the morphed agents; see Martini et al. [22].

Individual WMC were determined by calculating the Ospan, Sspan and Rspan scores separately and then averaging them for each participant in order to obtain a WMC composite score (see

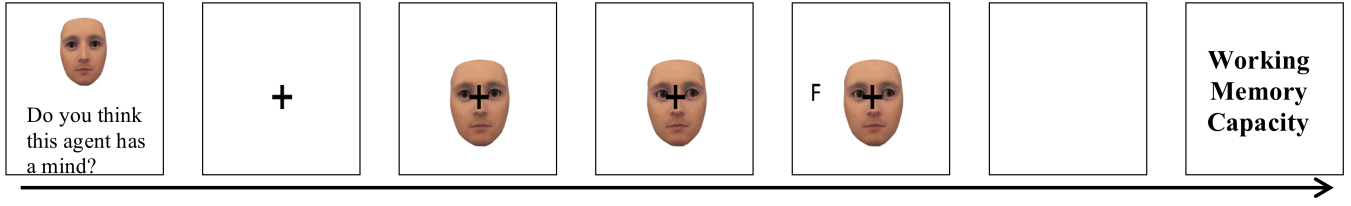


Figure 3: Sequence of events during the experiment: At the beginning of a block, participants rated the agent, then proceeded to gaze cueing trials: A trial began with a fixation cross, followed by appearance of agent face (60% agent shown in image), at which agent eyes shifted to either the right or to the left. After a target letter (T or F) appeared and remained on the screen until participant responded or for 1200 ms, whichever happened first. After gaze cueing trials, participants proceeded to the three measures of working memory

Table 1: Mean reaction times (and standard errors) during the gaze following task as a function of validity and humanness

Validity	0%	20%	40%	60%	80%	100%
Valid	469.85 (12.17)	473.60 (12.08)	480.80 (13.59)	471.38 (11.36)	477.49 (11.47)	478.45 (17.05)
Invalid	480.22 (12.13)	472.61 (11.69)	474.63 (11.65)	479.75 (13.06)	482.22 (12.74)	479.17 (13.44)

Kane et al., [17]). The WMC composite score was then correlated with individual difference scores per agent morph level. If individual WMC determined to what extent top-down control could be exerted on gaze following, we would expect a negative correlation between WMC and difference scores.

3 RESULTS

Reaction time data as a function of validity and % physical humanness are reported in Table 1. Difference scores as a function of individual WMC and % physical humanness are shown in Table 2, and correlations between individual WMC scores and difference scores are shown in Table 3.

As expected, the ANOVA of the mind rating data showed a significant main effect of % physical humanness ($F(1,36) = 11.21, p < .001$), with increasing mind ratings as a function of physical humanness. The analysis of the gaze following data revealed that participants were highly accurate in performing the target discrimination task ($M = 94.99\%$, $SE = .47\%$). The mean reaction time across all conditions was 476.70 ms ($SE = 11.32$ ms), with no significant difference in reaction times for valid compared to invalid trials ($p = .408$). The ANOVA of the difference scores revealed no main effect of % physical humanness ($F(1,36) = .81, p = .545$), indicating that the ability to exert top-down control did not vary as a function of agent appearance.

Correlational analyses between the WMC composite score and the difference scores revealed no significant negative relationship between WMC and difference scores in general, showing that WMC did not overall affect the extent of exerted top-down control ($r(37) = -.02, p = .80$). However, there was significantly stronger top-down control for participants with a high WMC than for participants with a low WMC capacity in response to the 80% human morph

Table 2: Difference scores between invalid and valid trials as a function WMC and physical humanness

WMC	0%	20%	40%	60%	80%	100%
High	8.54	1.90	-4.35	7.54	1.23	6.89
Low	21.12	-15.93	-15.58	12.71	22.82	-31.14

Table 3: Correlations between difference scores and WMC as a function of physical humanness

Variable	0%	20%	40%	60%	80%	100%
Pearson R	-.04	.19	.06	-.08	-.36*	.14
Significance	.801	.267	.739	.622	.029	.422

($r(37) = -.36, p = .029$). Correlations for none of the other agents were significant ($p > .267$).

4 DISCUSSION

The goal of the experiment was to examine if the ability to suppress reflexive shifts of attention in a counterpredictive gaze following paradigm (where targets appear with a higher likelihood at the uncued than the cued location) varies as a function of WMC, and in particular, if participants with a high WMC are better at suppressing reflexive shifts of attention in conditions where processing agent appearance is associated with additional workload (i.e., humanoid agents). The results show no differences in top-down control abilities between participants with high versus low WMC across morph levels. However, participants with high WMC were significantly better at suppressing reflexive shifts of attention in response to the 80% human morph, where categorical ambiguity is maximal [18]. Thus, when being cued by ambiguous humanoid agents, participants with high WMC were better able to inhibit shifts of attention to locations that most likely would not contain the target (i.e., cued location), and to execute shifts of attention to more likely target locations instead (i.e., uncued location).

The observation that participants with high WMC are better able to exert top-down control in more cognitively demanding conditions is in line with previous studies, stating that top-down control of gaze following depends on the context [6, 9, 20, 24, 31, 32, 34]. It is also in line with the categorical ambiguity explanation for the

Uncanny Valley, which states that agents of 70-80% humanness are hardest to categorize given that they are human-like enough to evoke perceptions of humanness, but not perfectly human, resulting in categorization ambiguity [18, 21, 22] and cognitive conflict [19]. With regard to the current data, this means that top-down control abilities might only be impaired when the gazer's physical appearance is so ambiguous that categorizing it as human versus non-human consumes significant amounts of cognitive resources, and that participants with high WMC are more able to make up for this significant loss of resources than participants with low WMC [19].

There were certain limitations to the experiment. First, the data is only preliminary and the small sample size might not provide enough power to detect individual differences in top-down control. In contrast to Martini et al. [20], the current study showed no reflexive orienting for all morphed agents but only for the 60% and 80% morphs, which might also be due to the small sample size. Although this shows that the most ambiguous agents (i.e., with the most ambiguous mind ratings) did impair top-down control, we are currently collecting more participants in order to attain a higher statistical power. Lastly, although following the gaze of the 80% morph was most strongly modulated by differences in WMC, it is not associated with the most positive difference score. This is probably due to the fact that the overall score constitutes the average score of strong reflexive orienting for participants with low WMC and strong top-down control for participants with high WMC.

Our results have implications for the field of human-robot interaction. Since robots with realistic, but not perfectly human physical appearance cause additional load on working memory, trying to increase anthropomorphism via physical appearance can have detrimental effects on performance in human-robot interaction, in particular if the task is demanding. Although the drainage of resources does not seem to be extreme in the current experiment, being exposed to agents with an ambiguous appearance could cause significant drainage of resources over time, which would particularly impair performance on resource-sensitive tasks like sustained attention. Future studies need to investigate the relationship between categorical ambiguity and cognitive performance in human-robot interaction further, with a particular focus on the design of physical agent features and their effect on the drainage of cognitive resources.

REFERENCES

- [1] Abdulaziz Abubshait and Eva Wiese. 2017. You Look Human, But Act Like a Machine: Agent Appearance and Behavior Modulate Different Aspects of Human-Robot Interaction. *Frontiers in Psychology* 8 (2017). <https://doi.org/10.3389/fpsyg.2017.01393>
- [2] Jeremy N. Bailenson, Andrew C. Beall, and Jim Blascovich. 2002. Gaze and task performance in shared virtual environments. *The Journal of Visualization and Computer Animation* 13, 5 (Dec. 2002), 313–320. <https://doi.org/10.1002/vis.297>
- [3] Andrew P. Bayliss, Alexandra Frischen, Mark J. Fenske, and Steven P. Tipper. 2007. Affective evaluations of objects are influenced by observed gaze direction and emotional expression. *Cognition* 104, 3 (Sept. 2007), 644–653. <https://doi.org/10.1016/j.cognition.2006.07.012>
- [4] Andrew P. Bayliss, Matthew A. Paul, Peter R. Cannon, and Steven P. Tipper. 2006. Gaze cuing and affective judgments of objects: I like what you look at. *Psychonomic Bulletin & Review* 13, 6 (Dec. 2006), 1061–1066.
- [5] Gary Bente, Sabine R  ijggenberg, Nicole C. Kr  dmer, and Felix Eschenburg. 2008. Avatar-Mediated Networking: Increasing Social Presence and Interpersonal Trust in Net-Based Collaborations. *Human Communication Research* 34, 2 (April 2008), 287–318. <https://doi.org/10.1111/j.1468-2958.2008.00322.x>
- [6] Robert O Deaner, Stephen V Shepherd, and Michael L Platt. 2007. Familiarity accentuates gaze cuing in women but not men. *Biology Letters* 3, 1 (Feb. 2007), 64–67. <https://doi.org/10.1098/rsbl.2006.0564>
- [7] N. J. Emery. 2000. The eyes have it: the neuroethology, function and evolution of social gaze. *Neuroscience and Biobehavioral Reviews* 24, 6 (Aug. 2000), 581–604.
- [8] Jeffrey L. Foster, Zach Shipstead, Tyler L. Harrison, Kenny L. Hicks, Thomas S. Redick, and Randall W. Engle. 2015. Shortened complex span tasks can reliably measure working memory capacity. *Memory & Cognition* 43, 2 (Feb. 2015), 226–236. <https://doi.org/10.3758/s13421-014-0461-7>
- [9] Elaine Fox, Andrew J. Calder, Andrew Mathews, and Jenny Yiend. 2007. Anxiety and Sensitivity to Gaze Direction in Emotionally Expressive Faces. *Emotion (Washington, D.C.)* 7, 3 (Aug. 2007), 478–486. <https://doi.org/10.1037/1528-3542.7.3.478>
- [10] Chris Kelland Friesen and Alan Kingstone. 1998. The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic bulletin & review* 5, 3 (1998), 490–495.
- [11] Chris Kelland Friesen, Jelena Ristic, and Alan Kingstone. 2004. Attentional Effects of Counterpredictive Gaze and Arrow Cues. *Journal of Experimental Psychology: Human Perception and Performance* 30, 2 (2004), 319–329. <https://doi.org/10.1037/0096-1523.30.2.319>
- [12] Alexandra Frischen, Andrew P. Bayliss, and Steven P. Tipper. 2007. Gaze Cueing of Attention. *Psychological bulletin* 133, 4 (July 2007), 694–724. <https://doi.org/10.1037/0033-2909.133.4.694>
- [13] Chris D. Frith and Uta Frith. 2006. The Neural Basis of Mentalizing. *Neuron* 50, 4 (May 2006), 531–534. <https://doi.org/10.1016/j.neuron.2006.05.001>
- [14] Kurt Gray, Liane Young, and Adam Waytz. 2012. Mind Perception Is the Essence of Morality. *Psychological Inquiry* 23, 2 (April 2012), 101–124. <https://doi.org/10.1080/1047840X.2012.651387>
- [15] Dana A. Hayward and Jelena Ristic. 2013. Measuring attention using the Posner cuing paradigm: the role of across and within trial target probabilities. *Frontiers in Human Neuroscience* 7 (May 2013). <https://doi.org/10.3389/fnhum.2013.00205>
- [16] Bruce M. Hood, J. Douglas Willen, and Jon Driver. 1998. Adult's eyes trigger shifts of visual attention in human infants. *Psychological Science* 9, 2 (1998), 131–134.
- [17] Michael J. Kane, M. Kathryn Bleckley, Andrew R. A. Conway, and Randall W. Engle. 2001. A controlled-attention view of working-memory capacity. *Journal of Experimental Psychology: General* 130, 2 (2001), 169–183. <https://doi.org/10.1037/0096-3445.130.2.169>
- [18] Jari K  dt  ryi, Klaus F  urger, Meeri M  dk  dr  dinen, and Tapio Takala. 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 (2015). <https://doi.org/10.3389/fpsyg.2015.00390>
- [19] Arielle R. Mandell, Melissa A. Smith, Molly C. Martini, Tyler H. Shaw, and Eva Wiese. 2015. Does the Presence of Social Agents Improve Cognitive Performance on a Vigilance Task? In *Social Robotics*, Adriana Tapus, Elisabeth Andr  AI, Jean-Claude Martin, Fran  ois Ferland, and Mehdi Ammi (Eds.). Vol. 9388. Springer International Publishing, Cham, 421–430. http://link.springer.com/10.1007/978-3-319-25554-5_42 DOI: 10.1007/978-3-319-25554-5_42
- [20] Molly C. Martini, George A. Buzzell, and Eva Wiese. 2015. Agent Appearance Modulates Mind Attribution and Social Attention in Human-Robot Interaction. In *Social Robotics (Lecture Notes in Computer Science)*. Springer, Cham, 431–439. https://doi.org/10.1007/978-3-319-25554-5_43
- [21] Maya B. Mathur and David B. Reichling. 2016. Navigating a social world with robot partners: A quantitative cartography of the Uncanny Valley. *Cognition* 146 (Jan. 2016), 22–32. <https://doi.org/10.1016/j.cognition.2015.09.008>
- [22] M. Mori, K. F. MacDorman, and N. Kageki. 2012. The Uncanny Valley [From the Field]. *IEEE Robotics Automation Magazine* 19, 2 (June 2012), 98–100. <https://doi.org/10.1109/MRA.2012.2192811>
- [23] Anna Pecchinenda and Manuel Petrucci. 2016. Emotion Unchained: Facial Expression Modulates Gaze Cueing under Cognitive Load. *PLOS ONE* 11, 12 (Dec. 2016), e0168111. <https://doi.org/10.1371/journal.pone.0168111>
- [24] Giuseppina Porciello, Marco Tullio Liuzza, Ilaria Minio-paluello, Gian Vittorio Caprara, and Salvatore Maria Aglioti. 2016. Fortunes and misfortunes of political leaders reflected in the eyes of their electors. *Experimental Brain Research; Heidelberg* 234, 3 (March 2016), 733–740. <https://doi.org/10.1007/s00221-015-4496-1>
- [25] Michael I. Posner. 1980. Orienting of attention. *Quarterly Journal of Experimental Psychology* 32, 1 (Feb. 1980), 3–25. <https://doi.org/10.1080/00335580808248231>
- [26] Susanne Quadflieg, Malia F. Mason, and C. Neil Macrae. 2004. The owl and the pussycat: Gaze cues and visuospatial orienting. *Psychonomic bulletin & review* 11, 5 (2004), 826–831.
- [27] Leonhard Schilbach, Marcus Wilms, Simon B. Eickhoff, Sandro Romanzetti, Ralf Tepest, Gary Bente, N. Jon Shah, Gereon R. Fink, and Kai Vogeley. 2009. Minds Made for Sharing: Initiating Joint Attention Recruits Reward-related Neurocircuitry. *Journal of Cognitive Neuroscience* 22, 12 (Nov. 2009), 2702–2715. <https://doi.org/10.1162/jocn.2009.21401>
- [28] Atsushi Senju and Gergely Csibra. 2008. Gaze Following in Human Infants Depends on Communicative Signals. *Current Biology* 18, 9 (May 2008), 668–671.

- <https://doi.org/10.1016/j.cub.2008.03.059>
- [29] Christoph Teufel, Dean M. Alexis, Nicola S. Clayton, and Greg Davis. 2010. Mental-state attribution drives rapid, reflexive gaze following. *Attention, Perception, & Psychophysics* 72, 3 (April 2010), 695–705. <https://doi.org/10.3758/APP.72.3.695>
 - [30] Eva Wiese, Tyler Shaw, Daniel Lofaro, and Carryl Baldwin. 2017. Designing Artificial Agents as Social Companions. *Proceedings of the Human Factors and Ergonomics Society Annual Meeting* 61, 1 (Sept. 2017), 1604–1608. <https://doi.org/10.1177/1541931213601764>
 - [31] Eva Wiese, Agnieszka Wykowska, and Hermann J. Mäijller. 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 (April 2014), e94529. <https://doi.org/10.1371/journal.pone.0094529>
 - [32] Eva Wiese, Agnieszka Wykowska, Jan Zwickel, and Hermann J. Mäijller. 2012. I See What You Mean: How Attentional Selection Is Shaped by Ascribing Intentions to Others. *PLOS ONE* 7, 9 (Sept. 2012), e45391. <https://doi.org/10.1371/journal.pone.0045391>
 - [33] Eva Wiese, Jan Zwickel, and Hermann Josef Mäijller. 2013. The importance of context information for the spatial specificity of gaze cueing. *Attention, Perception, & Psychophysics* 75, 5 (July 2013), 967–982. <https://doi.org/10.3758/s13414-013-0444-y>
 - [34] Agnieszka Wykowska, Eva Wiese, Aaron Prosser, and Hermann J. Mäijller. 2014. Beliefs about the Minds of Others Influence How We Process Sensory Information. *PLOS ONE* 9, 4 (April 2014), e94339. <https://doi.org/10.1371/journal.pone.0094339>