Predictive gaze cues affect face evaluations: The effect of facial emotion

Andrew P. Bayliss, Debra Griffiths, and Steven P. Tipper

School of Psychology, Bangor University, Bangor, UK

When we see someone change their direction of gaze, we spontaneously follow their eyes because we expect people to look at interesting objects. Bayliss and Tipper (2006) examined the consequences of observing this expectancy being either confirmed or violated by faces producing reliable or unreliable gaze cues. Participants viewed different faces that would consistently look at the target, or consistently look away from the target: The faces that consistently looked towards targets were subsequently chosen as being more trustworthy than the faces that consistently looked away from targets. The current work demonstrates that these gaze contingency effects are only detected when faces create a positive social context by smiling, but not in the negative context when all the faces held angry or neutral expressions. These data suggest that implicit processing of the reward contingencies associated with gaze cues relies on a positive emotional expression to maintain expectations of a favourable outcome of joint attention episodes.

Keywords: Eye gaze; Attention; Face processing; Social evaluation.

Responding appropriately to the social cues of other members of our social group is critical to establishing fluent interactions with other people. For example, if one sees someone suddenly move their eyes, the natural response is to establish “joint attention”, by orienting to the same object. Joint attention has been studied heavily in developmental psychology because of its critical role in social development (e.g., Baron-Cohen, 1995; Moore & Dunham, 1995); however, there has been an increased interest in the attentional mechanisms that underlie this orienting effect in adults (e.g.,

Correspondence should be addressed to Andrew Bayliss, School of Psychology, Bangor University, Brigantia Building, Penrallt Road, Bangor, Gwynedd LL57 2AS, UK. E-mail: a.bayliss@bangor.ac.uk

This work was supported by a Leverhulme Early Career Fellowship awarded to AB, and by a Wellcome Programme grant to ST. The MacBrain Face Stimulus Set was developed by Nim Tottenham (tott0006@tc.umn.edu), supported by the John D. and Catherine T. MacArthur Foundation.

Driver et al., 1999; Friesen & Kingstone, 1998; see Frischen, Bayliss, & Tipper, 2007, for review).

The advantage in following each others' gaze is because people usually attend to objects of interest in the environment, which we might not otherwise have become aware of in the ever-changing world. Hence, normally developing infants benefit greatly by following gaze spontaneously from a very young age (Scaife & Bruner, 1975), and this continues into adulthood. Indeed, gaze cues have effects on attention even when, in cognitive psychology paradigms, objects of interest (i.e., targets) only appear at the gazed location 50%, or even 25% of the time (Driver et al., 1999; Friesen, Ristic & Kingstone, 2004; Tipples, 2008). It appears therefore, in most circumstances, attention is cued by gaze direction in an automatic fashion.

The gaze following system developed to quickly orient to the direction of gaze because it is usually advantageous to do so. However, the very fluency of the system makes it vulnerable to deception. A "cheater", knowing that their social gaze will influence the attention of others, is able to gain an advantage by looking away from the locus of their planned action. For example, a skilled soccer player might look to the left, but then pass to the right, knowing that his opponent, despite often being comparably skilled, will anticipate a leftward pass. Magicians frequently use misdirection by using social cues in performances to remarkable effect (see Kuhn & Land, 2006). In both cases, the observer knows that deception is likely, but is unable to prevent being tricked.

People are nevertheless highly sensitive to the perceived trustworthiness of the people they meet. For example, manipulating the temporal dynamics of a smile can influence trustworthiness ratings of faces such that "fake" smiles cause reductions in perceived trustworthiness, despite other physical aspects of the smiles being kept constant (Krumhuber et al., 2007). Winston, Strange, O'Doherty, and Dolan (2002) showed that amygdala activation in response to a face is predicted by the trustworthiness of that face, even in a passive viewing task (see also Todorov, Gobbini, Evans, & Haxby, 2007). Similar effects have been shown in studies where trustworthiness is an attribute that is acquired via an experimental manipulation, rather than as an extant property of the face (e.g., Singer, Kiebel, Winston, Dolan, & Frith, 2004), where faces paired with uncooperative behaviour in economic exchange games are trusted less, and activate the amygdala more strongly than faces paired with cooperative behaviour.

Rather than using the outcomes of economic games to predict acquired trustworthiness, Bayliss and Tipper (2006) employed a gaze cueing paradigm using faces that behaved differently throughout the experiment. The participants were told that the faces would look towards and away from the targets with equal frequency. In fact, whenever certain faces appeared, they would *always* look at the target. Hence, these reliable faces helped

participants complete the task more efficiently, by causing them to shift attention to the target, thereby reducing reaction time. Furthermore, these faces also held to the social expectancy that people tend to look at objects of interest or importance. In contrast, intermixed with these faces, other identities always looked at the opposite location to where the target would appear. These counterpredictive faces hindered the participant, as they were consistently deceptive in their behaviour. Bayliss and Tipper found that the gaze direction of these predictive and counterpredictive faces produced identical cueing effects to a third set of faces that looked towards targets on 50% of trials, which is the standard proportion in such experiments. These cueing effects are somewhat surprising, as they suggest that the gaze cueing system appears to be insensitive to contingencies related to the reliability of the social cues produced by different people, and cannot modify the strength of the attention shift based on prior encounters (but see Deaner, Shepherd, & Platt, 2007; Frischen & Tipper, 2006, for evidence that gaze cueing can be influenced by person identity in some circumstances).

After the completion of the gaze cueing segment of the study, Bayliss and Tipper (2006) then presented participants with pairs of these faces and asked them to choose the most trustworthy face. Although participants were still unaware that one face from each pair had looked at the targets, and the other always looked away from the targets, participants chose the reliable "helpful" faces significantly more often than the "deceptive" faces. There was also evidence that the "deceptive" faces were better encoded in memory in their paradigm, since participants tended to feel that they had encountered these "deceptive" faces more often during the experiment than the "helpful" faces. This notion converges with evidence of enhanced memory for untrustworthy faces (Yamagishi, Tanida, Mashima, Shimoma, & Kanazawa, 2003). Interestingly, participants rarely noticed the gaze cueing contingencies of the faces, supporting the idea that these effects emerge through unconscious processes. It therefore appears that while gaze cueing itself cannot be calibrated to account for the likely reliability (or otherwise) of the social cues given by a known individual, the contingencies are nevertheless encoded and are able to affect subsequent evaluative judgments of the faces.

This study aims to replicate and extend the findings of Bayliss and Tipper (2006) by investigating the effect of facial emotion on the ability to detect deceptive individuals from gaze shift behaviour. In this experiment, the faces expressed different facial emotions to three separate groups of participants. Hence, each group of participants saw either happy, angry, or neutral facial expressions. There are two main motivations for this manipulation. The first is empirical: We intend to confirm and replicate the findings of Bayliss and Tipper. Investigating the effect further while controlling for and manipulating facial emotion has the potential of identifying important boundary

conditions for the emergence of the effect.¹ The second is theoretical: How efficient are we at picking up on the subtleties of social interactions in markedly different social contexts? Is the behaviour of others encoded in a similar way when every individual we encounter is in a positive state or negative emotional state? If this is the case, then we predict that the emotion of the faces producing gaze cues will have no effect on the strength to which the cue contingencies affect participants' judgements of the faces.

On the other hand, facial emotion may have a critical role. When someone smiles at you, a positive outcome to the social exchange is expected, whereas this is not expected when someone looks at you with an angry or neutral emotional expression. Might this shift in expectancy in these situations impact the way that social contingencies are picked up on in this paradigm? Indeed, the idea that learning is potentiated in situations where expectations are violated lends support to this hypothesis (Kamin, 1969; Rescorla & Wagner, 1972). That is, when a happy face violates the rule of looking at interesting objects, this will have a stronger impact on the system than when an angry face behaves in the same way. Hence, we predict stronger effects of gaze contingencies when the face is happy than when angry. An additional condition where neutral faces were used was included in the design, and the clear expectation is that the magnitude of the effects will fall in between those of the happy and angry conditions. That is, the positivity of the emotional expression will share a linear relationship with the measures of implicit learning of gaze cue contingencies.

METHODS

Participants

Seventy-two adults (20 males, 52 females), with a mean age of 20.3 years ($SD = 2.9$ years) participated in the experiment in return for course credit or payment. All were naïve as to the purpose of the experiment in general, and therefore naïve to the predictive nature of the faces' gaze cues. Informed consent was gained, and participants received course credit for their participation.

¹ The faces used in the original Bayliss and Tipper (2006) study were collected from various sources, and as such the photographs were rather heterogeneous in their composition. However, it is possible that the positive emotional expression on the majority of the faces used in the original study may have potentiated the effects on trustworthiness and memory.

Stimuli

Twenty unfamiliar faces were taken from a larger sample of 36 faces from the NimStim face database. These faces were arranged into pairs such that one face from each of the pairs belonged to Face Group A ($n = 10$), and the other to Face Group B ($n = 10$). The selection of the pairs was initially based on the “neutral expression” versions of the face stimuli. Each pair was matched for gender, ethnicity, and approximate age. Independent raters ($n = 12$) also ensured that the pairs of faces were rated equal attractiveness and trustworthiness, and that, as a whole, both groups of faces (A and B) were of approximately equal attractiveness and trustworthiness. There were two pairs of Black females, two pairs of Black males, three pairs of White females, and three pairs of White males. Examples of each face holding three emotional expressions, “happy”, “neutral”, and “angry”, were used. The eyes in the original stimuli were staring straight ahead. These stimuli were manipulated to produce two further versions of each face, with eyes looking to the left, and to the right. Hence there were a total of 180 face stimuli ($20 \text{ identities} \times 3 \text{ emotions} \times 3 \text{ gaze directions}$). The faces measured approximately $10.6 \times 10.0 \text{ cm}$. The eye region measured between 4.0 and 4.5 cm between the left corner of the left eye and the right corner of the right eye. The eyes themselves measured approximately $0.5 \times 1.0 \text{ cm}$, with pupils/irises measuring approximately $0.5 \times 0.5 \text{ cm}$.

Pictures of 36 household objects were used as target stimuli. Eighteen of these were categorised as belonging in the kitchen, and 18 were garage items. These could appear in one of four colours (red, green, blue, and yellow), and could appear flipped about the vertical axis (so that, for example, a kettle could appear with the handle on the left, or on the right). Targets varied $1.5\text{--}5.0 \text{ cm} \times 3.0\text{--}8.0 \text{ cm}$, and were presented such that their centre was 10.0 cm to the left or right of the centre of the screen. A black fixation cross was also used, measuring $0.7 \times 0.7 \text{ cm}$, placed at the centre of the screen. The stimuli were presented on a monitor placed approximately 60 cm from the participant.

Design

The factor “emotional expression” was manipulated between subjects, with participants split equally into groups of participants observing happy, neutral, or angry faces. The factor, “congruency” was manipulated within subjects. The eyes of the face could either look towards (“congruent”) or away from (“incongruent”) the target. For half the participants, the faces from Group A looked at the target consistently, whereas the faces from Group B consistently looked away from the target. The faces presented to the other half of the participants had the opposite condition assignment (Face Group A

looked away from, and Face Group B looked towards targets). With this procedure, the effect of any residual a priori differences between the trustworthiness and attractiveness of the pairs of faces was minimised.

Procedure

Gaze cueing. Participants were seated in front of the computer screen, with their heads placed on a chinrest. They were instructed to fixate the centre of the screen throughout each trial, and, while ignoring the ostensibly nonpredictive direction of gaze, to respond to the target as quickly and accurately as possible. Participants responded to “kitchen” items with the spacebar, and to “garage” items with the “h” key, using the thumb and first finger of their right hand, respectively. On each trial, a fixation cross appeared in the centre of the screen for 600 ms. This was followed by one of the 20 faces, staring straight ahead for 1500 ms. Then, 500 ms before the onset of the target, the eyes would move to the left or the right. Depending on whether the face was from the congruent or incongruent set of faces for a given participant, the target would then appear to the left or right of the screen. Half of the targets would be “garage”, and half “kitchen” items, although the exact identity, colour, and orientation of the target were selected randomly without replacement on each trial. At response (or after 2500 ms, whichever was sooner), the display cleared and a feedback tone sounded (bell or buzzer for correct or incorrect, respectively). This blank screen remained for 1500 ms before the next trial began (see Figure 1).

After completing 12 practice trials, in which a single face from a separate set would look at, and away from the target an equal number of times, participants completed a total of 240 trials, split over two blocks of 120. Hence, each face appeared 12 times over the course of the session. Which face was presented on each trial was selected randomly.

Face evaluations. Directly after the gaze cueing procedure, the participants completed three runs through a two-alternative forced choice procedure in which they were shown displays of the ten pairs of faces (i.e., one from Group A, and one from Group B). These were presented next to each other (to the left and right of fixation). The order of the pairs was randomised, as was the left–right positioning of each face within the pairs. For each face pair, participants were asked to choose which face they felt was the most trustworthy of the pair. They made this response by hitting “1” on the keypad if they felt the face on the left was most trustworthy, and “2” if the face on the right was the most trustworthy of the pair. A blank screen was presented after each response for 2000 ms.

After choosing the most trustworthy face from each of the 10 pairs, participants completed the questionnaire a second time (face pair order and

Figure 1. (A) Examples of a “congruent” and an “incongruent” gaze cue trial. (B) Examples of face pairs for each emotion condition. The stimuli were presented in colour. Contact the authors for a colour version of this figure, or view the online version of the journal.

face position in a new random order). This time, participants were told that some faces might have appeared in the experiment more often than others (in fact, they had all appeared 12 times). Hence, they were asked to choose from each pair the face they felt had been presented more often during the gaze cueing experiment.

Next, participants were debriefed fully. This involved asking whether they had noticed the fact that some faces had acted differently to others. No subject noticed anything close to the true experimental manipulation. Then, they were told about the differing contingencies of the two sets of faces. It was ensured that each participant understood the manipulation fully. That is, that each of the face pairs they viewed, one face had always looked at the target, and that the other face always looked away from the target.

Manipulation check. Although participants were unable to explicitly recall the gaze contingencies to which they had been exposed, it was important to examine this in a more formal test. Hence participants were then asked to repeat the questionnaire a final time, but this time, they were

asked to choose the face that they thought might have been the face that had consistently looked at the target throughout the gaze cueing experiment. Participants were indeed at chance at this task, correctly identifying the congruently cueing faces on only 49.6% of trials, $t(71) = -0.23$, $p = .8$, (50.4%, 52.5%, and 45.8% for happy, neutral, and angry groups, respectively, all nonsignificant). Hence, we can be confident that any effects of gaze contingency on the central measures described in this report are not due to the implementation of conscious strategies on the part of the participants: They clearly had no conscious awareness of the gaze cue contingencies.

RESULTS

Gaze cueing

Target categorisation errors were made on 6.4% of incongruent and 5.9% of congruent trials. A mixed-factor ANOVA with “cueing” as a within-subjects factor (congruent vs. incongruent) and “emotion” as a between-subjects factor (happy, neutral, and angry) revealed no significant effects, all F s < 1.6 , p s $> .2$. Along with trials with erroneous responses, trials with RTs > 1500 ms (2.5%) were removed prior to analysis of mean RT. A similar analysis to that performed on errors was conducted, revealing a significant main effect of cueing, $F(1, 69) = 21.5$, $MSE = 290$, $p < .001$, $\eta_p^2 = .24$ because RTs were quicker to targets appearing in congruently cued locations (719 ms) than targets appearing at incongruently cued locations (733 ms). The remaining main effect and interaction were both nonsignificant (F s < 1). Nevertheless, t -tests were performed on the RT data for each group to ensure each emotion elicited gaze cueing reliably. Significant gaze cueing was observed in the happy face group, $t(23) = 2.87$, $p = .009$, $d = 0.59$, where congruent (699 ms) were faster than incongruently cued targets (715 ms); the neutral face group, $t(23) = 2.94$, $p = .007$, $d = 0.60$, congruent (722 ms) faster than incongruent targets (735 ms); and finally, the angry faces, $t(23) = 2.24$, $p = .035$, $d = 0.46$, congruent (737 ms) faster than to incongruent targets (748 ms).

Facial evaluations

When asked to choose which was the most trustworthy from between a pair of faces, participants chose the face that had consistently looked at the target throughout the gaze cueing experiment an average of 5.56 times out of 10 pairs of faces (see Figure 2). A one-sample t -test demonstrated that this tendency was above chance, $t(71) = 3.38$, $p = .001$, one-tailed, $d = 0.40$. On the second run through the forced-choice procedure, participants were requested to choose the face that they felt they had seen more often during the gaze

Figure 2. Mean number of “congruent” faces chosen for each measure, for each experiment. The y-axis is intersected around chance (5/10). Error bars represent standard error of the mean. Asterisks indicate p -values derived from one-sample t -test contrasts; *** $p < .001$, * $p < .05$, + $p < .10$.

cueing experiment, based on the false premise that one may have appeared more often than the other. Here, participants chose the face that had consistently gazed to the target an average of 4.71 times out of 10 pairs. A one-sample t -test demonstrated that this face frequency response was significantly below chance (i.e., 5 out of 10 pairs), $t(71) = -1.71$, $p = .045$, one-tailed, $d = 0.20$. That is, participants more often felt that the incongruently gazing faces had appeared more often during the gaze cueing experiment.

These data essentially replicate and confirm those of Bayliss and Tipper (2006). When faces are seen to consistently predict the location of a visual target by always looking towards it they are evaluated as being more trustworthy in comparison to faces that consistently deceive the participants by looking away from the target. Importantly, the participants chose these “helpful” faces significantly *less* often in the second task, where they were considering frequency of exposure by choosing the face they felt had appeared more often.

The effect of facial emotion

As outlined in the introduction to this paper, we hypothesised that a positive emotion expressed by the faces in this experiment will potentiate the effects

of predictive gaze on face evaluation. Thus, the specific prediction is that as a face increases in positive valence, the ratings of trustworthiness and frequency of exposure will diverge from chance; that the proportion of congruent gazing faces chosen as more trustworthy will increase, but fewer will be chosen as having appeared more often in the experiment than the counterpredictive faces. Hence, a Measure $\times$ Emotion ANOVA was performed, followed by planned contrasts for each measure, testing for differences between the happy, angry, and neutral groups. This ANOVA revealed a significant effect of measure, $F(1, 69) = 13.1$, $MSE = 1.97$, $p = .001$, $\eta_p^2 = .16$. This demonstrates again that responses on the two measures deviate from each other since fewer predictive-valid faces were chosen in the frequency measure as compared with the trustworthiness measure. The main effect of emotion was not significant, $F(1, 69) < 1$. The critical Measure $\times$ Emotion interaction was marginal, $F(1, 69) = 2.93$, $MSE = 1.97$, $p = .06$, $\eta_p^2 = .078$. The planned contrasts revealed that this trend was due to the participants in the happy group showing a significantly stronger effect of trustworthiness than participants viewing angry faces, $p = .05$, one-tailed, and a trend for a stronger frequency effect, $p = .055$, one-tailed. No other contrasts approached significance.

To confirm the effects of face emotion on the gaze cueing contingency effect on trust and face frequency, separate analysis of each emotion group was undertaken (see Figure 2). As expected, the effect for trust, $t(23) = 4.61$, $p < .001$, $d = 0.94$, and frequency of presentation, $t(23) = -2.10$, $p = .047$, $d = 0.43$, were significant when faces smiled. However, for the neutral faces the trust effect was marginal, $t = 1.81$, $p = .083$, $d = 0.37$, and there was no significant appearance frequency effect, $t(23) = -0.50$, *ns*. Finally for the faces expressing anger there were no significant effects, $ts(23) = 0.71$ and -0.15 , *ns*.

GENERAL DISCUSSION

This study aimed to confirm and extend the results of Bayliss and Tipper (2006), which showed that faces that reliably looked at targets were subsequently judged to be more trustworthy than faces that looked away from targets. These gaze cue contingencies also had effects on memory: Episodes with deceptive faces appeared to be better encoded in memory leading to increased familiarity with these faces. In the current study, stimuli that were additionally controlled for emotional expression were used. The central findings were that the effects on social judgements and memory were only reliably replicated when the faces used were smiling. Despite the fact that angry and neutral faces produced reliable gaze cueing effects, the

contingencies between an individual face and their gaze behaviour failed to affect its perceived trustworthiness.

In addition to the replication of the effects on trustworthiness and frequency measures of face memory, we felt it important to investigate the possibility that the effects were due in part to explicit memory of the gaze contingencies associated with the different faces. Participants did not report that they had noticed that some faces were looking at the target each time they appeared, and others always looked away, so a more formal measure was employed. Interestingly, after being informed of the gaze contingencies participants were completely unable to correctly identify the faces that had reliably indicated the target location. This suggests that where the effect does emerge, it might not be due to conscious processing of the observed social behaviour of the faces. Furthermore, in contrast to Bayliss and Tipper (2006), we did not use any “filler” faces with random gaze target contingencies. Hence, even where *all* faces were consistent in their behaviour, participants still failed to notice the experimental manipulation.

Of central importance, the magnitude of the effects of gaze contingencies on the trustworthiness and memory measures did appear to be linearly related to the valence of the facial emotion. That is, the effects of gaze contingencies increased as the emotional expression of the faces became more positive (see Figure 2). This suggests that these effects are indeed modulated by emotional expression, rather than any lower level structural features of the stimuli used. This may be because when we encounter angry faces, the expectation of a positive social exchange is quite low, so when this weak expectation is violated by the eyes of the face looking away from the target, little impact is made on the cognitive system and hence only a weak link between the facial identity of the angry face and its behaviour is possible. The converse is true with a happy face where we have a high expectation of a positive outcome, and hence a stronger impact, in terms of learning, is made when this social expectation is violated (Kamin, 1969; Rescorla & Wagner, 1972).

A further interpretation of our results can be made from a mood induction perspective. Bolte, Goschke, and Kuhl (2003) found that participants' ability to discriminate sets of associated words from sets of unassociated words was significantly better following the induction of a positive mood as compared with a negative mood induction procedure and baseline. Importantly, participants were no more able to consciously solve these word-based problems, suggesting that these effects emerged through processes that are implicit and affected by emotional context. This is related to our finding since if being exposed to happy faces improves mood, then according to Bolte et al. this could promote “the activation of widespread associative networks” (p. 420). Such activation would potentially facilitate connections between the observed social behaviour of individuals and the

representation of their facial identity, thus boosting the effect on trustworthiness and memory observed here in the happy group of participants.

Finally, we have to acknowledge memory processes that might mediate the effect. For example, it is possible that the gaze target contingencies are better encoded when faces are smiling because when the negative outcome occurs, as in the case of an incongruent gaze cue, the impact on the evaluation of that face would be greater as it mismatches with the positive smiling context. Furthermore, given that long-term memory appears to be better for positive emotional faces (Shimamura, Ross, & Bennett, 2006), the effect of prior gaze contingencies leaves a stronger trace in memory when the face is smiling, influencing later assessments of individuals for trust. Of course all these concepts, learning via expectancy violation, mood-modulated association generation and memory for emotional faces are all intertwined and are all likely to have an impact in the present paradigm and contribute to the impact of facial emotion on the effects we describe.

In conclusion, participants in this study were unable to prevent their attention orienting in the direction of another person's gaze, even when it was consistently deceptive about the location of a target. This is perhaps because the neural networks responsible for gaze perception and identity processing are separate, and hence interactions between gaze and identity may only emerge later in the processing stream (e.g., Andrews & Ewbank, 2004; Hoffman & Haxby, 2000). Subtle interactions between gaze direction and person identity nevertheless do indeed take place, influencing judgements of trust, but at a level not available for conscious report. That gaze cueing contingencies influence judgements of trust again demonstrate the sophistication of social evaluative systems, and their ability to guide decision making in the social environment. Although we might not be able to prevent orienting in the direction in which a known deceiver is looking, the deception is at least encoded, and this information can aid the guidance of other, more complex social interactions with such individuals.

Original manuscript received February 2008

Revised manuscript received October 2008

First published online January 2009

REFERENCES

- Andrews, T. J., & Ewbank, M. P. (2004). Distinct representations for facial identity and changeable aspects of faces in the human temporal lobe. *NeuroImage*, 23, 905–913.
- Baron-Cohen, S. (1995). The eye direction detector (EDD) and the shared attention mechanism (SAM): Two cases for evolutionary psychology. In C. Moore & P. J. Dunham (Eds.), *Joint attention: Its origins and role in development* (pp. 41–59). Hove, UK: Psychology Press.
- Bayliss, A. P., & Tipper, S. P. (2006). Predictive gaze cues and personality judgments: Should eye trust you? *Psychological Science*, 17(6), 514–520.

- Bolte, A., Goschke, T., & Kuhl, J. (2003). Emotion and intuition: Effects of positive and negative mood on implicit judgments of semantic coherence. *Psychological Science*, 14(5), 416–421.
- Deaner, R. O., Shepherd, S. V., & Platt, M. L. (2007). Familiarity accentuates gaze cueing in women but not men. *Biology Letters*, 3(1), 63–67.
- Driver, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual Cognition*, 6, 509–540.
- Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin and Review*, 5, 490–495.
- Friesen, C. K., Ristic, J., & Kingstone, A. (2004). Attentional effects of counterpredictive gaze and arrow cues. *Journal of Experimental Psychology: Human Perception and Performance*, 30, 319–329.
- Frischen, A., Bayliss, A. P., & Tipper, S. P. (2007). Gaze-cueing of attention: Visual attention, social cognition and individual differences. *Psychological Bulletin*, 133(4), 694–724.
- Frischen, A., & Tipper, S. P. (2006). Long-term gaze cueing effects: Evidence for retrieval of prior states of attention from memory. *Visual Cognition*, 14, 351–364.
- Hoffman, E. A., & Haxby, J. V. (2000). Distinct representations of eye gaze and identity in the distributed human neural system for face perception. *Nature Neuroscience*, 3, 80–84.
- Kamin, L. J. (1969). Predictability, surprise, attention, and conditioning. In B. A. Campbell & R. M. Church (Eds.), *Punishment and aversive behavior* (pp. 279–296). New York: Appleton-Century-Crofts.
- Krumhuber, E., Manstead, A. S. R., Cosker, D., Marshall, D., Rosin, P. L., & Kappas, A. (2007). Facial dynamics as indicators of trustworthiness and cooperative behavior. *Emotion*, 7(4), 730–735.
- Kuhn, G., & Land, M. F. (2006). There's more to magic than meets the eye! *Current Biology*, 16(22), R950.
- Moore, C., & Dunham, P. J. (Eds.). (1995). *Joint attention: Its origins and role in development*. Hove, UK: Psychology Press.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. H. Black & W. F. Prokasy (Eds.), *Classical conditioning II: Current research and theory* (pp. 64–100). New York: Appleton-Century-Crofts.
- Scaife, M., & Bruner, J. S. (1975). The capacity for joint visual attention in the infant. *Nature*, 253(5489), 265–266.
- Shimamura, A. P., Ross, J. G., & Bennett, H. D. (2006). Memory for facial expressions: The power of a smile. *Psychonomic Bulletin and Review*, 13(2), 217–222.
- Singer, T., Kiebel, S. J., Winston, J. S., Dolan, R. J., & Frith, C. D. (2004). Brain responses to the acquired moral status of faces. *Neuron*, 41, 653–662.
- Tipples, J. (2008). Orienting to counterpredictive gaze and arrow cues. *Perception and Psychophysics*, 70(1), 72–87.
- Todorov, A., Gobbini, M. I., Evans, K. K., & Haxby, J. V. (2007). Spontaneous retrieval of affective person knowledge in face perception. *Neuropsychologia*, 45, 163–173.
- Winston, J. S., Strange, B. A., O'Doherty, J., & Dolan, R. J. (2002). Automatic and intentional brain responses during evaluation of trustworthiness of faces. *Nature Neuroscience*, 5(3), 277–283.
- Yamagishi, T., Tanida, S., Mashima, R., Shimoma, E., & Kanazawa, S. (2003). You can judge a book by its cover: Evidence that cheaters may look different from cooperators. *Evolution and Human Behavior*, 24, 290–301.