Seeing Seeing

Ben Phillips

(Forthcoming in *Philosophy and Phenomenological Research*)

Abstract: I argue that we can visually perceive others as seeing agents. I start by characterizing perceptual processes as those that are causally controlled by proximal stimuli. I then distinguish between various forms of visual perspective-taking, before presenting evidence that most of them come in perceptual varieties. In doing so, I clarify and defend the view that some forms of visual perspective-taking are “automatic”—a view that has been marshalled in support of dual-process accounts of mindreading.

1. Introduction

Suppose you are walking towards the front door of your house. You see an overly chatty neighbor and immediately break out into a jog. As you pass through the doorway, you glance back and notice that you have been spotted. Here is the question: did you *visually perceive* your neighbor as seeing you, or, did you visually perceive his eyes and acquire the post-perceptual belief that he sees you?

Several philosophers have converged on the view that we can perceive basic emotions and motor intentions (e.g., Gallagher, 2008a, 2008b; Scholl & Gao, 2013; Gallagher & Varga, 2014; Helton, 2018; Varga, 2018). However, little attention has been paid to the question of whether we can perceive visual states. This is somewhat surprising, for there is a tight connection between seeing and gazing behavior. If your eyes are pointing in my direction, that is a straightforward visual cue that you see me. It follows that if there are heuristics, built into the visual system, that guide attributions of seeing, they could be very relatively simple ones. For example, it is possible that the visual system deploys a basic heuristic of the following form: *If X’s eyes point in the direction of O then X sees O.*

In what follows, I argue that some forms of visual perspective-taking are indeed perceptual. I start by specifying what I mean by “perception,” elaborating on the widely-held view that perceptual processes are distinctive in being causally controlled by proximal stimuli (sect. 2). In section 3, I review research on gaze perception. I then outline the different forms of visual perspective-taking that gaze perception supports. In section 4, I examine those forms of visual perspective-taking that require gaze-following. I present evidence that these forms of perspective-taking are perceptual. I also defend and clarify the view that they are “automatic.” In section 5, I shift the focus to those forms of visual perspective-taking that do not require gaze-following. For instance, I focus on attributions of the form *THAT AGENT SEES ME* and *THAT AGENT IS LOOKING FEARFULLY AT SOMETHING*. I argue that the visual system produces both kinds of attributions.¹

2. The perception/cognition distinction

Various theorists have claimed that we can perceive emotions and motor intentions. But in making these sorts of claims, how are these theorists demarcating the perceptual from the cognitive?

2.1. Stimulus control

According to the direct social perception thesis (DSP), we sometimes “directly perceive” the mental states of others. Proponents of DSP often cash out the term “direct perception” in terms of the distinction between inferential and non-inferential processes. For instance, Gallagher (2008a) likens the perception of a mental state to the non-inferential perception of a car:

¹ I will use the following terms as synonyms for “visual perspective-taking”: “representing someone as seeing something”; “representing a visual perspective”; and “attributing a state of seeing.”

At the personal or conscious level, I do not have to perceptually piece together the shape and the color and the mass in order to get my car. Even if the sub-personal processes are complex (and I do not deny that they are), the perception that I have of my car is direct—I see it right there in front of me. I do not have to glue anything together, add an interpretation or add an inference. (2008a, 537)

Gallagher (2008a, 2008b) goes on to argue that some personal-level representations of mental states are similarly direct.

According to a view that is widely held in cognitive science, all perceptual processes are inferential (e.g., see Pylyshyn, 1999; Palmer, 1999; Frisby & Stone, 2010). The onus is thus on the DSP-theorist to provide a notion of inference that (i) does some explanatory work, and (ii) does not apply to paradigmatically perceptual processes (see Gallagher & Vargas, 2014).

In what follows, I'll sidestep these issues by adopting an alternative approach to the perceivability question. The approach I'll adopt is based on the simple observation that perceptual processes are causally controlled by proximal stimuli in a way that post-perceptual processes are not.² For instance, in discussing whether we can visually perceive animacy, Scholl and Gao (2013) work with the following notion:

[A] hallmark feature of perception (vs. cognition) is its strict dependence on subtle visual display details; percepts seem to be irresistibly controlled by the nuances of the visual input regardless of our knowledge, intentions, or decisions. (2013, 209)

² I am not denying that one can draw the perception/cognition distinction in terms of the distinction between inferential and non-inferential processes. It is unlikely that there is only one explanatorily fruitful way to mark a border between perception and cognition (Beck, 2018; Phillips, 2019a).

Other proponents of the view that perceptual processes are distinctive in being “controlled” by proximal stimuli include Palmer (1999, 5); Kanwisher (2001, 90); Frisby and Stone (2010, 16); Herschbach (2015); and Beck (2018).

What does it mean to say that a process is “controlled” by proximal stimuli? Consider some examples from other domains. Genes are often construed as causally controlling the sequencing of amino acids (Godfrey-Smith, 2000; Stegmann, 2014); the movement of a driver’s limbs causally controls the trajectory of the car; and the movement of the disc and needle causally controls the sequence of sounds emitted from the record player. What unifies these processes is that they are all what Stegmann (2014) calls “external orderings”: a sequence of events, such that each member was caused by an event that was not part of the sequence itself.

We can use this notion of causal control to characterize perception as follows: a process is perceptual, not cognitive, just in case it has the function of representing the distal causes of proximal stimuli, where these representations are produced in a way that is causally controlled by these stimuli. Call any process that has this function, “stimulus-controlled.”³

One key virtue of this account is that it accommodates token cognitive states that were caused by a proximal stimulus. For example, suppose I look at some deciduous trees, an experience which causes me to acquire the post-perceptual thought, *THERE WON’T BE ANY LEAVES ON THE TREES NEXT MONTH*. This thought does not qualify as perceptual because it does not have the *function* of being causally-controlled by proximal stimuli. It is a thought that I can maintain as I watch television or lie in

³ What determines the function of a process? According to *etiological accounts*, the function of a trait is determined by its contribution to past reproductive success (e.g., see Godfrey-Smith, 1994). According to *propensity accounts*, the function of a trait is determined by its potential contribution to future reproductive success (e.g., see Bigelow & Pargetter, 1987). Other non-etiological accounts are defended by Cummins (1975) and Nanay (2010). In claiming that perceptual processes have the function specified above, I intend to remain neutral on this debate.

bed at night: both episodes during which the stream of proximal stimuli fails to causally control my sequence of thoughts.

Some processes are reliably triggered by proximal stimuli, but fail to qualify as perceptual because they do not have the function of representing the distal causes of these stimuli. Semantic priming is a case in point. Just seeing a patch of red can automatically activate representations of associated words, such as “tomato” and “blood” (Nijboer et al., 2006). But we would not construe this as a distinctively *perceptual* process. Rather, priming is commonly construed as a process of *spreading activation*: the prime activates a network of semantically-related representations, which has been forged through a process of associative learning (see Neely, 1991). Spreading activation is a domain-general process. Any node can serve as the prime, triggering the activation of other representations in the network: the prime needn’t be a proximal stimulus. But even when the prime *is* a proximal stimulus, the process of spreading activation does not have the function of representing the distal cause of this stimulus: its function is to simply activate other representations in the network.

In sum, then, to qualify as perceptual, a process must have the dual function of (i) being causally controlled by proximal stimuli, and (ii) representing the distal causes of these stimuli.

2.2. Stimulus control and automaticity

I have just characterized perceptual processes as stimulus-controlled processes. Before determining whether any visual perspective-taking processes are stimulus controlled, it will be useful to address a closely related way of characterizing perceptual processes.

According to some theorists, perceptual processes are automatic in a way that cognitive process are not (e.g., see Fodor, 1983; Pylyshyn, 1999; Scholl & Gao, 2013).

For example, Scholl and Gao (2013) couple their characterization of perception as stimulus-controlled with the following remark:

In general, “perception” refers in this context to a family of processes that are relatively automatic and irresistible ... (2013, 202)

What is it for a process to be automatic? Within cognitive science, “automatic” is not a univocal term (Moors & De Houwer, 2006). Some theorists use it to denote processes that are unintentional, involuntary, or uncontrollable; others use it to denote processes that are unconscious or pre-attentive. For our purposes, it will suffice to focus on two kinds of process that are often run together: the mandatory and the ballistic.

Let’s say that a process, *P*, is *mandatory* just in case it meets the following criterion: once the system housing *P* receives an input from its proper domain, *P* begins regardless of which other mental processes are occurring within the given individual (Fodor, 1983; Mandelbaum, 2014). Let’s say that a process is *ballistic* just in case, once triggered, it cannot be stopped via endogenous psychological means (e.g. due to a shift in goals). In contrast to mandatory processes, ballistic processes do not necessarily begin once the proper input has been received; however, unlike mandatory processes, once a ballistic process starts, it cannot be stopped via endogenous psychological means (Logan & Cowan, 1984; Bargh, 1992, 186; Mandelbaum, 2014).

By way of illustration, consider color perception. Once you attend to an object, your visual system goes about its business, determining the object’s color. Moreover, once color processing has begun, you cannot just decide to perceive the given object achromatically. Color perception is thus a good candidate for a process that is both mandatory and ballistic.

How is stimulus control related to mandatoriness and ballisticsity? Notice that for a process to be mandatory just is for that process to be causally controlled by certain

domain-specific inputs. If these domain-specific inputs are proximal stimuli then the process qualifies as stimulus-controlled, and thus perceptual. Any evidence for a mandatory process that takes visible gazing-behavior as input and produces representations of seeing as output will therefore constitute evidence for *perceptual* visual perspective-taking.

What if we uncover a merely ballistic process that takes visible gazing-behavior as input and produces representations of seeing as output? Will this process qualify as stimulus-controlled and thus perceptual?

To answer this question, let's start with an intuition pump. You pay a neuroscientist to install a device that suppresses the activity of color-processing neurons when you want to see the world in black-and-white. The device can pre-empt the activity of color-processing neurons, but it cannot suppress their activity once triggered. After the operation, color-processing thus becomes a ballistic, but not a mandatory, affair for you. Are your representations of color still *perceptual*?

It seems to me that you have retained your ability to see colors. The reason is that when you decide to see the world in color, you have no control over *which* colors you see objects as being: the stimuli are still in control. For instance, if you come across a banana which, due to lighting conditions, looks blue, you cannot just decide to see it as yellow. You still succumb to the illusion. Your representation of the banana's color thus retains a key hallmark of perception: namely, insensitivity to countervailing evidence. Why has it retained this hallmark? Because your color-processing mechanisms have retained the function of representing objects' colors in a stimulus-controlled manner.

The procedure you underwent only altered *when* color-processing occurs, not how it unfolds once triggered.⁴

To sum up: any mandatory or ballistic process that takes visible gazing-behavior as input and produces representations of seeing as output will qualify as perceptual. This distinction between mandatory and ballistic forms of visual perspective-taking will be important below, for as we will see, the main argument that visual perspective-taking is not “automatic” only targets the thesis that it is mandatory, not the thesis that it is ballistic.

3. Gaze perception

In assessing whether we can see others’ visual perspectives, a natural starting point is gaze perception. The strongest evidence that we visually perceive others’ gaze comes from studies of selective adaptation to gaze *direction*.

3.1. Perceiving gaze direction

Selective adaptation is a paradigmatically perceptual phenomenon whereby a neuron’s firing rate drops amidst prolonged exposure to its preferred stimulus. This decrease in firing rate raises the threshold for detecting the feature in question, while lowering the threshold for detecting novel features. For example, in the waterfall illusion, staring at the downward motion of the water fatigues those neurons responsible for detecting downward motion. The result is that if one stares at a stationary object immediately afterwards, it will look like it is moving upwards.

Adaptation is a ubiquitous feature of perceptual processes. It is not, however, a feature of paradigmatically cognitive states. For example, merely thinking about a

⁴ I am not claiming that being insensitive to countervailing evidence is sufficient for being perceptual. Rather, I am claiming that being stimulus-controlled *explains* why perceptual processes tend to be insensitive to countervailing evidence.

waterfall does not induce thoughts of upward motion. For these reasons, adaptation is widely used as a marker of the perceptual (e.g., see Block, 2014). For our purposes, it is a useful marker because it provides strong evidence of stimulus control. If the outputs of a system exhibit adaptation effects when the system receives certain stimuli as inputs, this is good evidence that the system’s outputs are causally controlled by those stimuli.

Various studies have found evidence of visual adaptation to gaze direction (Jenkins et al., 2006; Calder et al., 2008; Teufel et al., 2009; Bayliss et al., 2011). In the study by Jenkins et al. (2006) subjects adapted to an agent’s gaze, which was directed 5 to 10 degrees off center. Subjects then immediately saw a second agent whose direction of gaze was within the same range. Jenkins and colleagues found that subjects inaccurately saw the second agent as looking directly at them, even though the agent’s gaze was averted.

By varying low-level features of the stimulus, researchers have been able to rule out the hypothesis that subjects are merely adapting to low-level features, such as shapes and colors (Bayliss et al., 2011). The neurophysiological evidence supports this, with various studies converging on the view that different neurons in the superior temporal sulcus code for different directions of gaze (Perrett et al., 1992; De Souza et al., 2005; Calder et al., 2007).

3.2. From gaze perception to visual perspective-taking

The evidence thus suggests that we visually perceive gaze direction, but for genuine visual perspective-taking to occur, the outputs of gaze-perception systems must feed into systems that produce representations of seeing. Are any of *these* systems perceptual? If so, what sorts of seeing-attributions do they produce?

In the literature, most researchers focus on a distinction between two levels of visual perspective-taking: Level 1, which is the representation of *which* objects someone else

sees; and Level 2, which is the representation of *how* those objects are seen (Flavell, 1977).

According to a view which has dominated discussions of visual perspective-taking, only Level 1 is automatic (Apperly & Butterfill, 2009; Apperly, 2011; Butterfill & Apperly, 2013). I examine the evidence below. However, before doing so, let's start with a more fine-grained taxonomy. Consider the following attributions:

- (1) *S* sees that as yellow.
- (2) *S* sees that.
- (3) *S* sees me.
- (4) *S* is looking angrily/fearfully/happily at something in that direction.

(1) is an instance of Level 2 perspective-taking, while (2) is the paradigmatic form of Level 1 perspective-taking. (3) and (4) are instances of Level 1 perspective-taking; however, they are importantly different from (2) in that they do not require one to first perceive *S*'s direction of gaze, and then shift one's attention to the object at which *S* is looking. In other words, unlike (2), gaze-following is not required for (3) and (4).

Given that these different forms of visual perspective-taking will likely recruit different mechanisms—mechanisms that address different ecological needs—it is entirely possible that some are perceptual, while others are exclusively cognitive. Thus, in what follows, rather than focusing on the coarse-grained Level 1/Level 2 distinction, I review the evidence concerning each form of perspective-taking listed above.

4. Paradigmatic Level 1 perspective-taking

In the literature on automatic perspective-taking, the focus has been on representations of the form, *S SEES THAT*—the paradigmatic form of Level 1 perspective-taking.

Given that attributions of this sort require one to perceive the given agent's direction of

gaze, before locating any objects in that direction, two experimental paradigms have dominated the literature: gaze-cueing experiments; and the dot perspective task.

4.1. Gaze-cueing and the dot perspective task

In standard gaze-cueing experiments, subjects are tasked with identifying a letter, which appears either to the left or to the right of a centrally-located face. Even when subjects are told that the face is irrelevant, they cannot help but follow its gaze. This is evidenced by the fact that subjects are faster at locating and identifying the target letter when the face gazes towards it. When the face gazes in the opposite direction, subjects are slower and less accurate (Friesen & Kingstone, 1998; Driver et al., 1999).

Some theorists have taken this as evidence that gaze-following is an automatic stimulus-driven process (e.g., Friesen & Kingstone, 1998). Even so, gaze-following is not tantamount to visual perspective-taking. To follow your gaze, all I have to do is shift my attention in the direction that your eyes are pointing: I need not represent you as *seeing* anything. This brings us to the dot perspective task.

In the dot perspective task, subjects view a scene containing an avatar, along with two dots. In the self-perspective condition, subjects report how many dots they see in the room. In the other-perspective condition, they report how many dots the avatar can see. When the avatar’s perspective differs from their own (e.g. the avatar can only see one of the dots), subjects’ responses are slower and more error prone. Importantly, this occurs not just when subjects are asked to report on the avatar’s perspective, but also when they are asked to report on their own (Samson et al., 2010). Subjects thus appear to compute the avatar’s visual perspective, even when they know that it is irrelevant to the task at hand. This is referred to as an “altercentric effect.”

Additional evidence comes from a study by Qureshi and colleagues (2010), in which subjects were asked to make the judgments outlined above while performing a task that

taxed executive function.⁵ Interestingly, this secondary task did not attenuate the effects outlined above: in fact, it enhanced them. This suggests that the process of adopting the avatar’s perspective is not carried out in central cognition; rather, it is an involuntary stimulus-controlled process.

Various theorists have taken these findings as evidence for automatic visual perspective-taking: we do not just automatically detect and follow the gaze of others; we automatically represent the visual perspectives that lie behind their gazing behavior (Apperly & Butterfill, 2009; Samson et al., 2010; Apperly, 2011; Butterfill & Apperly, 2013; Capozzi et al., 2014; Nielson et al., 2015; Baker et al., 2016; Furlanetto et al., 2016).

For some theorists—most notably, Apperly and Butterfill—the dot perspective task provides evidence that there are two systems for visual perspective-taking: an early developing, phylogenetically ancient, system which carries out Level 1 perspective-taking in an automatic fashion; and a late-developing, phylogenetically young, system which carries out Level 2 perspective-taking in a slow and deliberative manner. The motivation for denying that Level 2 perspective-taking is automatic comes from variants of the original dot perspective task. For instance, Surtees et al. (2010) set up a scenario in which subjects viewed the numeral “9,” which looked like a “6” from the avatar’s vantage point. This difference in *how* the object was seen by subject and avatar did not generate an altercentric effect.⁶

⁵ Executive function encompasses the set of cognitive processes responsible for flexible goal-directed behavior (e.g. planning, working memory, response inhibition, and set-shifting).

⁶ In fact, when it comes to the question of automaticity, studies of Level 2 perspective-taking have generated mixed results. See Westra (2017) for a useful discussion. In what follows, I set Level 2 perspective-taking aside. The strongest candidates for automatic perspective-taking are instances of Level 1.

4.2 *Three challenges to the automatic-perspective-taking hypothesis*

For our purposes, the findings outlined above suggest that attributions of the form *S SEES THAT* are computed in a stimulus-controlled, and thus perceptual, manner. However, this conclusion faces three main challenges. According to one challenge, gaze-following is controlled by cognitive states, such as beliefs and expectations, and therefore does not qualify as automatic. According to another challenge, there is no compelling evidence for a system dedicated to Level 1 perspective-taking; rather, the gaze-cueing effects described above are controlled by domain-general mechanisms of attention. According to the third challenge, even if there is a system dedicated to automatic gaze-following, the evidence suggests that it does not output representations of *seeing*. Let's consider each challenge in turn.

4.2.1. *Is gaze-following automatic?*

Teufel and colleagues (2010) utilized a version of the standard gaze-cueing task. However, instead of a static image of a face, subjects were shown a video of an experimenter wearing a pair of goggles. When the goggles were transparent, subjects exhibited reflexive gaze-following. However, when the goggles were opaque, and subjects were thereby led to believe that the experimenter could not see, there was a significant reduction in reflexive gaze-following.

Subsequent studies have uncovered similar effects. For instance, Wiese and colleagues (2014) varied subjects' beliefs about whether a robot's gaze was controlled by a human. Subjects reflexively followed the robot's gaze only when they believed that it was controlled by a human. Perez-Osorio and colleagues (2015) varied whether an agent's direction of gaze was congruent with her intention to act. Importantly, subjects were made aware of these intentions during the pre-cueing phase. Perez-Osorio and colleagues

found that subjects were less likely to exhibit reflexive gaze-following when the agent gazed towards an action-incongruent object.⁷

What are we to make of these top-down effects on gaze-following? At the very least, they show that neither gaze-following nor the forms of perspective-taking that rely on gaze-following are mandatory. Does this mean that visual perspective-taking is voluntary and post-perceptual? Not if we agree that perceptual processes can be merely ballistic.

Recall that ballistic processes are not necessarily triggered by their proper inputs, but once triggered, other mental processes cannot stop them from running to completion. The suite of top-down factors that modulate gaze-following all modulate whether it is triggered. They appear to play this triggering role by modulating the degree to which the subject attends to the agent's gaze: a minimum threshold of attention being required for the gaze-following process to begin.

For instance, recall the experiment carried out by Teufel and colleagues (2010), in which they varied subjects' beliefs about whether a goggle-wearing agent could see. Teufel and colleagues found that reflexive gaze-following did not occur if the subject believed that the agent could not see. Importantly, though, Teufel and colleagues (2009, 2013) also tested whether these beliefs modulate adaptation to gaze direction. What they found is that if the subject believes that the agent can see, adaptation to gaze direction occurs (e.g. prolonged exposure to leftward gaze causes subjects to see subsequent gazes as shifted to the right). In contrast, they found that if the subject believes that the agent cannot see, adaptation to gaze direction either does not occur (Teufel et al., 2009) or is significantly attenuated (Teufel et al., 2013).

⁷ Other studies exhibiting similar effects include Nuku and Bekkering (2008); Teufel et al. (2009, 2010); Ricciardelli et al. (2013); Wykowska et al. (2014); Dalmasso et al. (2016); Terrizzi and Beier (2016); and Morgan et al. (2018).

These findings strongly suggest that one’s belief regarding whether another agent can see modulates the degree to which one attends to that agent’s gaze. It appears that a minimum threshold of attention is required for reflexive gaze-following to begin, with top-down factors, such as beliefs and expectations, affecting whether this threshold is met: if I believe that you cannot see, I won’t be motivated to attend to your eyes as much. However, once initiated, the gaze-following process seems to elude voluntary control.

If this picture is correct, it means that gaze-following is a ballistic process the function of which is to track gaze in a way that is controlled by the stimulus. If top-down factors can influence whether the process *begins*, this does not alter the fact that it performs this function.⁸

4.2.2. *Is there a system dedicated to gaze-following?*

The second challenge to the claim that representations of the form *S SEES THAT* are produced automatically is based on studies in which gaze-cueing effects are generated using stimuli other than faces. For example, some studies have found a gaze-cueing effect when swapping out faces for arrows (Tipples, 2002; Kuhn & Kingstone, 2009). Santiesteban and colleagues (2013) altered the dot perspective task, replacing the avatar with an arrow of similar color, size, and shape. They found that regardless of whether it

⁸ Further studies may well show that cognitive factors affect *how* gaze-following unfolds, not just whether it is triggered. Would this undermine the view that some forms of gaze-following are stimulus-controlled? Not necessarily, for suppose it turns out that some gaze-following processes do not *require* these cognitive factors in order to run to completion. These processes would arguably qualify as stimulus-controlled, even though they can be influenced by cognition. The same could be said about paradigmatically perceptual processes, such as color-processing. Suppose, as some have claimed, cognitive states can penetrate color-processing mechanisms, distorting one’s perceptions of color (e.g., see Hansen et al., 2006). Color-processing presumably retains its status as stimulus-controlled, despite the potential influence of these cognitive factors, for it does not require them in order to run to completion. In other words, these cognitive factors do not have a constitutive role to play in guiding color-processing: they are what we might call, *contingent influences*. Thank you to an anonymous referee for raising this issue.

was an arrow pointing to one dot or an avatar gazing at the same dot, subjects' performances were largely the same: that is, the arrow/avatar interfered with their judgments concerning how many dots they could see.

Given that an arrow can have the same attention-cueing effects as an agent's gaze, some have concluded that there is no system dedicated to gaze-following; rather, instances of gaze-following are controlled by domain-general mechanisms of attention. For example, Santiesteban and colleagues (2013) suggest that it is the "directional asymmetry" of the avatar's profile that causes subjects to gaze in the same direction as it. Heyes (2014) calls this appeal to domain-general attentional processes the "submentalizing" hypothesis.

The submentalizing hypothesis is a potential death-knell for the view that representations of the form, *S SEES THAT*, are produced automatically. According to proponents of the submentalizing hypothesis, there is no mechanism dedicated to gaze-following, let alone visual perspective-taking. We only reflexively follow the gaze of others due to non-social cues, such as asymmetries in shape or color. These shifts in attention may eventually issue in attributions of seeing, but they are not driven by processes that have the function of representing seeing.

At this point, an obvious reply is to concede that eyes and arrows have comparable effects on attention, but to claim that distinct mechanisms are responsible in each case. For instance, it is plausible that through a process of domain-general associative learning, we have become habituated to looking where arrows point, whereas, the gaze-cueing effect is driven by a specialized gaze-following mechanism.

The evidence reviewed above does not adjudicate between the separate-mechanisms hypothesis and the single-mechanism hypothesis. But should we favor the latter on grounds of parsimony? This appeal to parsimony would have traction were there no

independent evidence for a mechanism dedicated to gaze-following. But there is plenty of such evidence.

First, as explained above, there is converging evidence for a system dedicated to the perception of gaze direction. Moreover, the gaze-cueing effect appears to be modulated by adaptation to gaze direction (Bayliss et al., 2011; Teufel et al., 2009, 2013). Second, not all asymmetries in low-level features trigger automatic shifts of attention. Ristic and Kingstone (2004) showed subjects an image that was ambiguous between a car and a face. The object’s directional asymmetry was thus held constant, but subjects only exhibited involuntary shifts of attention when they saw the image as containing a pair of eyes. Finally, arrows and eyes produce different cueing effects. Arrows cause subjects to reflexively shift their attention to any object on the congruent side, whereas, eyes cause subjects to shift their attention to the specific location where they point (Marotta et al., 2012; Gardner et al., 2018).

Overall, then, the evidence suggests that we cannot assimilate gaze- and arrow-cueing to the same, domain-general, processes of attention.

4.2.3. Does the gaze-following system produce representations of seeing?

The final challenge concerns the outputs of the gaze-following system. Even if the perception of eye-direction triggers an automatic process of gaze-following, why think that this process produces representations of *seeing*? Perhaps all the gaze-following system does is produce a shift in attention, congruent with the agent’s direction of gaze.

To test this deflationary hypothesis, Cole and colleagues (2015) altered the standard gaze-cueing protocol by adding a condition in which the agent’s line-of-sight to the target was blocked by a barrier. Despite the barrier’s presence, subjects still exhibited a gaze-cueing effect, comparable to the one found when the barrier was absent. Cole and colleagues (2016) also carried out a variant of the dot perspective task, in which the

avatar’s view of the given dot was blocked by a barrier. An altercentric effect was found, even though the avatar could not see the dot.⁹

In each study, evidence for reflexive gaze-following was found, even though the agent being observed did not have visual access to the target object. According to Cole and colleagues (2016), this casts serious doubt on the claim that reflexive gaze-following is a form of visual perspective-taking. Perhaps the mechanism responsible for reflexive gaze-following only produces representations of eye-direction: a geometrical feature that falls short of *seeing*.

Other studies have produced different results. Baker et al. (2016) found that putting barriers in between avatars and target objects did in fact remove the altercentric effect, while Freundlieb et al. (2017) achieved similar results using blindfolds (see also Furlanetto et al., 2016).

Given these mixed findings, additional studies are clearly needed. But suppose these studies converge on the view that the mechanisms of automatic gaze-following are insensitive to the effect that opaque barriers have on visual access. Would it follow that these mechanisms do not produce representations of *seeing*? Answering this question requires us to address an important question concerning the lower bounds of visual perspective-taking.

4.2.4. *The lower bounds of visual perspective-taking*

The proponent of the challenge we are currently considering harbors an assumption that often figures in debates about animal mindreading. Members of various species can distinguish between agents with visual access to an object, and those whose visual access is occluded by a barrier. In the most famous study, Hare and colleagues (2000) found

⁹ Wilson et al. (2017) got similar results using a blindfolded avatar. See also Conway et al. (2016), Cole et al. (2017), and Kuhn et al. (2018).

that when in the presence of a dominant chimpanzee, subordinate chimpanzees will only retrieve a piece of food if the dominant's line-of-sight to it is blocked by an opaque barrier: if the barrier is transparent, subordinates will not attempt to retrieve the food. For many theorists, this constitutes evidence that chimpanzees are Level 1 visual perspective-takers (see Buckner, 2014).

But what if chimpanzees had repeatedly failed this sort of task? To qualify as a visual perspective-taker, is it necessary to understand the way in which opaque barriers block visual access? This question becomes pressing once we realize that to represent seeing, an individual need not represent all its features, let alone accurately (see Buckner, 2014). To insist otherwise would be to adopt an unpalatable form of descriptivism about mental content (see Stich, 1996).

As folk psychologists, even human adults fail to represent seeing in an exhaustive and completely accurate fashion. For example, the folk are arguably wrong about the richness of visual experiences (Schwitzgebel, 2008). There are also features that the folk do not represent at all, such as the distinction between dorsal and ventral vision (Milner & Goodale, 1995). Despite these gaps and inaccuracies, we do not conclude that the folk fail to represent seeing.

Consider a more extreme case: a deluded man believes that people can see through opaque objects (e.g. he doesn't bother closing his curtains for the sake of privacy). Just because this man's beliefs about seeing are seriously mistaken, it does not follow that he lacks the ability to represent others' visual perspectives. But then why think that if *the visual system* is insensitive to the effects that barriers have on visual access, it does not represent seeing?

The hidden assumption at work appears to be this. Even if the deluded man described above does not understand the effect that barriers have on visual access, presumably he has *some* accurate beliefs about seeing. For example, perhaps he

understands that one can see the world inaccurately; that seeing leads to believing; and so on. These other folk-psychological beliefs are sufficient for the man to qualify as having a concept of seeing. If, however, *the visual system* is insensitive to the way in which opaque barriers remove visual access, it is unlikely that it represents these other aspects of seeing. Thus, at best, the visual system represents gaze direction—or so the argument goes.

But is it true that the visual system does not represent any other features of seeing? If in fact it does, it will be akin to our deluded man, who is a visual perspective taker, despite harboring mistaken views about the relation between opaque barriers and visual access.

One key feature of seeing concerns its causal relations to other mental states, such as emotions and motor intentions. Some theorists have argued that if an individual reliably tracks some of these causal relations, she qualifies as possessing *a* conception of seeing, however minimalistic (Wellman, 2014; Butterfill & Apperly, 2013; Phillips, 2019b). For example, Wellman argues that infants are more than gaze-followers: they understand that “seeing often produces inner, subjective experiences such as emotions and desires as well” (2014, 86).

We can think of an individual who understands at least some of these causal relations as possessing a functional conception of seeing. For instance, if a creature understands that gazing at *O* is inducing fear in *S*, this creature is more than a gaze-follower: it possesses a basic functional conception of seeing.

To clarify: I am not claiming that functionalism is the correct theory of seeing. Rather, I am making the more modest claim that to possess *a* conception of seeing, and to thereby qualify as a visual perspective taker, it suffices to represent at least some of the causal roles distinctive of seeing. If an individual fails to represent any of these roles then, at best, this individual is a mere gaze-follower.

Our question has therefore become this: Is *the visual system* a mere gaze-follower, or, does it represent at least some of the causal connections between gazing at an object and mental states, such as emotions and motor intentions?

It would certainly make adaptive sense for the visual system to represent some of the causal roles distinctive of seeing. If someone takes on a fearful expression, this signals danger somewhere; but only in combination with gaze direction does this cue indicate the precise location of the danger, as well as the direction in which you should run. Similarly, if someone takes on an angry expression, this signals danger to someone; but only in combination with gaze direction does it indicate whether *I* am the one facing the imminent threat (Adams & Kveraga, 2015).

Below, I present evidence that the visual system integrates cues for seeing and cues for other mental states in a stimulus-controlled manner. If what I just argued is correct, this provides support for the view that we visually perceive *seeing*, not just gaze direction.

4.2.5. *The wolf pack effect*

The first piece of evidence comes from a series of experiments by Gao and colleagues (2010). In one experiment, they found that the visual perception of chasing is modulated by perceived gaze direction. Subjects moved a disc (“the sheep”) around an array of moving darts. In the *Don’t Get Caught Task*, subjects were to avoid capture by “the wolf” (a bright red disc). In the *Wolf Pack Condition*, the moving distractor-darts always “faced” the sheep, while in the other condition, they always faced the wolf (Figure 1).

Figure 1. Recreated screen displays from the *Don't Get Caught Task* (adapted from Gao, McCarthy, & Scholl, 2010, 1850). (a) is an example of the *Wolf Pack Condition*, in which distractors face the sheep; while (b) is an example of the condition in which the distractors faced the wolf.

Even though subjects knew that the darts were irrelevant to the task, they were up to 20% less successful in the *Wolf Pack Condition* (Figure 1(a)), compared to a control in which the darts were orientated 90 degrees away from the sheep. This effect was greatly attenuated in the other condition, during which the darts faced toward the wolf, not the sheep (Figure 1(b)). According to Gao and colleagues, the best explanation for these findings is that in the *Wolf Pack Condition*, subjects misperceived the darts' random motion as intentional chasing. This misperception thereby distracted them from the task at hand: namely, avoiding the wolf. Why *misperception* as opposed to misjudgment? In addressing this question, Gao and colleague's reason as follows:

[T]hese results seem impossible to explain by appeal to explicit higher-level decisions, because what they show is how impotent such decisions can be. After all, subjects in

this experiment had every reason to simply ignore the wolfpack (and not to treat it as animate at all) since it was irrelevant or even disruptive to their overt task. But, despite these powerful task demands, they could not simply decide not to treat the wolfpack as animate. And this limitation is especially noteworthy here, since subjects were fully informed prior to the experiment about the irrelevance of the darts' orientations. These sorts of results strongly suggest that the resulting data reflect some properties and constraints of automatic perceptual processing rather than higher-level decisions that subjects are overtly making about the contents of the displays. (2010, 216)

In the *Leave Me Alone Task*, subjects had to move the sheep around with the aim of not touching any distractor-darts. The array was divided up into four quadrants: two contained wolf packs (i.e. darts that always faced the sheep), while the other two contained darts that always faced away from the sheep. Subjects spent considerably less time in the wolf pack quadrants. Using the same reasoning as above, Gao and colleagues (2010) conclude that the perception of intentional pursuit is processed in a stimulus-controlled, and thus perceptual, manner.

For our purposes, the key finding is that in each of the experiments described above, the perception of pursuit was modulated by the perception of gaze direction: more specifically, whether the pursuer was gazing at the pursued. Subjects knew that the direction of the darts was irrelevant to the task at hand. Regardless, when the darts faced the sheep, subjects could not help but represent them as pursuing the sheep. The best explanation for this finding is that cues for motor intentions and cues for the direction of visual attention are being integrated in a stimulus-controlled manner: evidence for *integration* comes from the fact that gaze-following cues modulate the impression of pursuit; while evidence for *stimulus control* comes from the fact that the

impression of pursuit cannot be suppressed in a top-down manner (e.g. by task goals or the knowledge that the darts are just mindless shapes).¹⁰

One might object that the darts did not have eyes, and it is therefore misleading to construe subjects as perceiving the darts' direction of visual attention. However, we frequently use body orientation as a cue for visual attention. In fact, cell recordings in macaques have revealed that a significant portion of neurons in the superior temporal sulcus respond not just to eyes, but to heads, and to bodies as well: these neurons are activated so long as the eyes/head/body points in the given direction (Perrett et al., 1992). The assimilation of all these disparate cues strongly suggests that the neurons in question process direction of visual attention.

Moreover, in one experiment, Gao and colleagues (2010) used discs with eyes, instead of darts (Figure 2). In the *Leave Me Alone Task*, they replicated the results obtained using darts: that is, subjects avoided those quadrants in which the discs' eyes pointed toward the sheep (see also van Buren & Scholl, 2016).¹¹

¹⁰ Helton (2018) and Varga (2018) both provide compelling arguments that subjects visually perceive the darts *as pursuing* the sheep. In making her case, Helton (2018) appeals to the fact that the impression of pursuit is, in an important sense, “unrevisable”—something she takes to be characteristic of perception, but not cognition. Varga (2018) appeals to the “irresistible and mandatory” nature of the impression of pursuit. I do not have the space to elaborate on these arguments. Suffice it to say, when we mark a perception/cognition border in terms of stimulus control, unrevisability and irresistibility become useful indicators of a perceptual process.

¹¹ You can experience the wolf pack effect for yourself at <http://perception.yale.edu/Brian/demos/Animacy-Wolfpack.html>

Figure 2. A recreated example of the wolf pack effect, generated using discs with eyes, as opposed to darts (adapted from Gao, McCarthy, & Scholl, 2010, 1848).

Together, these findings provide evidence that the visual system is not just sensitive to gaze direction: it is sensitive to the social significance of gaze, using it as a cue for the recognition of intentional action. Thus, even if the visual system is not sensitive to the effects that opaque barriers have on visual access, it tracks other features of seeing over and above gaze direction.

4.2.6. The automatic processing of gaze and affect

Additional evidence comes from studies of the automatic integration of cues for gaze direction and cues for emotional states. Mumenthaler and colleagues (2015) tasked subjects with recognizing dynamic expressions of emotion on a centrally-located target face. A “contextual face” appeared simultaneously in the periphery and was immediately backward masked (to forestall additional processing). Importantly, the contextual face only appeared for 30 ms. For the first 10 ms it had a neutral expression; then, during the middle 10 ms, it shifted its gaze either towards the target face or away from it. In

the final 10 ms, the contextual face exhibited an emotional expression (e.g. anger), before being backward masked. During this same 30 ms period, the target face went from a neutral expression to an expression of fear or anger, which coincided with the contextual face’s expression of emotion in the final 10 ms.

Mumenthaler and colleagues found that subjects recognized fear on the target’s face more easily when the contextual face looked at it angrily. In contrast, when the contextual face either looked away from the target face with an angry expression or looked at the target with a fearful expression, subjects did not recognize the target’s fear as easily.

Given that the contextual face’s expression appeared so rapidly (10 ms) and was backward masked, Mumenthaler and colleagues conclude that subjects automatically processed the “functional relationship between the contextual angry face and the target fearful face” (2015, 397). In other words, subjects’ visual systems automatically integrated cues for seeing and cues for emotion.

This conclusion is further supported by the fact that the earliest feedback loops within vision are estimated to take at least 50 ms to complete their round trip: top-down influences from central cognition take much longer (see Mandelbaum, 2018). There is thus not enough time for one’s beliefs about the contextual face to influence one’s identification of the target’s emotional expression. The identification is more likely to be the product of a stimulus-controlled process.

We therefore have compelling evidence that the visual system is sensitive to some of the causal-functional relations between seeing and emotion. The visual system appears to automatically produce representations of the form, *THAT AGENT IS LOOKING FEARFULLY/ANGRILY AT THIS AGENT*: a variety of Level 1 perspective-taking that is decidedly more sophisticated than the mere perception of gaze direction (e.g. *THOSE EYES ARE POINTING IN THAT DIRECTION*).

5. Automatic perspective-taking without gaze-following

I have just argued that the visual system produces representations of the form *S SEES THAT* and *S IS LOOKING FEARFULLY/ANGRILY AT T*. But recall that Level 1 perspective-taking comes in other forms. We should also consider representations such as *S SEES ME*; *S SEES SOMETHING IN THAT DIRECTION*; and *S IS LOOKING FEARFULLY/ANGRILY AT SOMETHING IN THAT DIRECTION*.

Importantly, the latter seeing-attributions do not require any gaze-following on the attributor's part. Gaze-following is a two-stage process during which one perceives another individual's direction of gaze, and then shifts one's own attention in the same direction. If what I have argued above is right, this second stage sometimes results in automatic seeing-attributions. But suppose you reject this claim. It does not follow that these other forms of visual perspective-taking are not automatic either, for none of them require the gaze-following stage. For instance, recognizing that you see *me* does not require that I shift my attention away from your eyes to some third party or object; neither does recognizing that you are looking fearfully at something in *that* direction (e.g. over my shoulder). These forms of perspective-taking must be considered on their own terms.

5.1. The automatic discrimination of direct versus averted gaze

Various studies have examined the perception of direct versus averted gaze. In a series of experiments, Adams and colleagues have gathered evidence for the view that we automatically detect averted gaze, so long as it is coupled with a fearful expression; and direct gaze, so long as it is coupled with an angry expression.

In one experiment, Adams and colleagues (2011) presented subjects with faces exhibiting either anger or fear, and either direct or averted gaze. Neural activity was recorded using fMRI. In the first experiment, each face was present for 2 seconds: long

enough for the subject to process it in a conscious and deliberative manner. In the second experiment, each face was presented for 33 ms, after which a neutral face (gazing in the same direction) was presented for 150 ms. This second face served as a backward mask, thereby ensuring that the target face was only processed in a stimulus-controlled manner (recall that 33 ms is too swift for the engagement of central cognition).

Adams and colleagues found that when subjects had enough time to process target faces in a conscious and deliberative manner, angry ones yielded the greatest activity when their gaze was averted, and fearful ones yielded the greatest activity when their gaze was direct. This activity occurred in the left amygdala, superior temporal sulcus, and other socio-cognitive areas. On the other hand, when subjects did not have enough time to process the faces in a reflective manner, the results were reversed: angry faces yielded the greatest activity when gaze was direct, and fearful faces yielded the most activity when gaze was averted. Interestingly, areas related to motor planning, such as the premotor cortex and the supplementary motor area, were more active in this condition.

What explains this pattern of results? According to Adams and colleagues, the key difference is between clear and ambiguous threat cues. If someone looks at you with an angry expression, that is a clear indication that you are in danger. The same goes for the case in which someone looks over your shoulder with a fearful expression (e.g. at the hungry leopard behind you). In contrast, suppose someone either looks at you with a fearful expression or looks over your shoulder with an angry expression. Neither is a clear indication that you are in danger: these cues are ambiguous, requiring further interpretation. According to Adams and colleagues, this is why socio-cognitive areas process clear threat cues in a rapid and reflexive manner, whereas, they process ambiguous threat cues in a slow and reflective manner. If I face a clear and imminent threat, I need to respond rapidly.

As Adams and colleagues point out, this interpretation of the evidence is buttressed by the fact that threatening stimuli appear to be processed separately along two neural pathways. The “low road” is a direct subcortical pathway proceeding from the thalamus to the amygdala: it appears to control automatic and rapid responses to threatening stimuli. The “high road,” however, recruits both cortical and subcortical networks: it is slow and deliberative, drawing on background information in a flexible manner. Importantly, there is evidence that the “low road” proceeds through areas of the amygdala which process both gaze direction and emotional expression (LeDoux, 1996; Burra et al., 2013; Méndez-Bortolo et al., 2016; Adams et al., 2017).

For our purposes, the evidence gathered by Adams and colleagues suggests that cues for seeing and cues for emotion are integrated automatically. If you are exposed to an angry face, looking directly at you, these findings suggest that you *visually represent* the face as such.

Behavioral studies lend further support to this view. For instance, some studies have shown that people are faster at identifying a face’s expression when its direction of gaze constitutes a clear threat cue (Adams & Kleck, 2003; Milders et al., 2011). Other studies have found that people’s interpretations of ambiguous or neutral expressions are biased according to the direction of gaze: the bias is towards seeing faces with averted gaze as expressing fear, and faces with direct gaze as expressing anger (Adams & Kleck, 2003, 2005). Finally, some studies have found that judgments of gaze direction are modulated by facial expression. For example, there is a bias in favor of seeing angry faces as exhibiting direct gaze, even if their gaze is somewhat averted (Graham & LaBar, 2007; Lobmaier, 2008; Ewbank et al., 2009; Rhodes et al., 2012).

One of these studies is worth elaborating on. Milders and colleagues (2011) got participants to perform an attentional blink task. The attentional blink occurs when an observer fails to identify a given target because it is embedded in a rapid stream of

visual stimuli, appearing shortly after another target—the observer’s attention is thought to be sapped by the first target. However, if an attentional blink does not occur (i.e. the observer is able to detect the given target), this constitutes evidence that targets belonging to the relevant category are processed in a stimulus-controlled, pre-attentive, manner.

In the study by Milders and colleagues, the targets were two faces. The first was a neutral face, and the second exhibited either an angry, fearful, or happy expression. In some trials, the emotional face looked directly at the subject; while in others, its gaze was averted. Both targets appeared for 80 ms, before being backward masked by a scrambled face. The subject’s task was to identify the gender of the first face (by pressing one of two keys), and to simply specify whether they saw a second face (also by pressing a key).

Milders and colleagues found that the detection of fearful, angry, and happy faces was modulated by gaze direction. If gaze was averted, there was an advantage in detecting fearful faces over angry or happy ones. However, if gaze was direct, this pattern was reversed. Given that target faces only appeared for 80 ms and were backward masked, Milders and colleagues draw the following conclusion:

What these results suggest is that the selection for awareness was not based on a single facial feature (e.g., wide open eyes). Instead, emotional expression and gaze direction of the face were both included in the evaluation process to determine the intention and motivation of the other person and the relevance to oneself, and this interaction of expression and gaze information appears to occur rapidly and automatically. (2011, 1460)

This automatic integration of cues for seeing and cues for emotion provide further evidence that we do not merely perceive eye-direction: we perceive others as looking at us angrily, looking away from us fearfully, and so on.¹²

6. Conclusion

In determining whether visual perspective-taking is ever perceptual, we need to go beyond the coarse-grained Level 1/Level 2 distinction. To that end, I have examined the following kinds of representations: *S SEES THAT; S IS LOOKING FEARFULLY/ANGRILY AT T; S SEES ME*; and *S IS LOOKING ANGRILY/FEARFULLY AT SOMETHING IN THAT DIRECTION*. There is no a priori reason to think that each type of representation will fall on the same side of the perception/cognition border. Nonetheless, I have argued that the visual system does in fact produce representations of each form. In making my case, I presented evidence that the visual system integrates cues for seeing and cues for other mental states, such as emotions and motor intentions.

The weight of evidence thus suggests that the visual system is keyed to *seeing*, not just direction of gaze. One way of putting this is to say that the visual system harbors a basic functional conception of seeing. It may not represent every causal relation between seeing, other mental states, and behavior; nonetheless, it represents at least some. The visual system is thus akin to our fallible folk psychologist, who represents others as seeing agents, even though he does not represent seeing in an exhaustive and completely accurate fashion.

¹² The mechanism responsible for the automatic integration of these cues may be a domain-specific adaptation. However, it is also possible that automatic integration occurs due to perceptual learning. More specifically, the visual integration of cues for emotion and gaze direction is a possible example of what Goldstone (1998) calls “unitization.” Unitization occurs when someone starts perceiving a single property, where previously they had perceived distinct properties. It may be that, through perceptual learning, I come to see others as gazing-angrily-at-me, gazing-fearfully-in-that-direction, and so on.

In assessing whether visual perspective-taking is automatic and perceptual, the moral is that researchers should look beyond the standard gaze-cueing paradigms. Not all forms of visual perspective-taking require gaze-following. Thus, even if gaze-following turns out to be non-automatic, it will not follow that visual perspective-taking is non-automatic across the board. Other experimental paradigms that interrogate different forms of visual perspective-taking need to be examined before general conclusions can be drawn.¹³

References

- Adams, R. B. Jr., Albohn, D. N., & Kveraga, K. (2017). Social vision: Applying a social-functional approach to face and expression perception. *Current Directions in Psychological Science*, 26(3), 243–248.
- Adams, R. B. Jr., Franklin, R. G., Jr., Nelson, A. J., Gordon, H. L., Kleck, R. E., Whalen, P. J., & Ambady, N. (2011). Differentially tuned responses to severely restricted versus prolonged awareness of threat: A preliminary fMRI investigation. *Brain and Cognition*, 77, 113–119.
- Adams, R. B. Jr., & Kleck, R. E. (2003). Perceived gaze direction and the processing of facial displays of emotion. *Psychological Science*, 14, 644–647.
- Adams, R. B. Jr., & Kleck, R. E. (2005). Effects of direct and averted gaze on the perception of facially communicated emotion. *Emotion*, 5, 3–11.
- Adams, R. B. Jr., & Kveraga, K. (2015). Social vision: Functional forecasting and the integration of compound cues. *Review of Philosophy and Psychology*, 6, 591–610.
- Apperly, I. A. (2011). *Mindreaders: The cognitive basis of “theory of mind.”* Hove, England: Psychology Press.
- Apperly, I. A., & Butterfill, S. A. (2009). Do humans have two systems to track beliefs and belief-like states? *Psychological Review*, 116(4), 953–970.
- Baker, L.J., Levin, D.T., & Saylor, M.M. (2016). The extent of default visual perspective taking in complex layouts. *Journal of Experimental Psychology: Human Perception and Performance*, 42, 508–516.
- Bargh, J. (1992). The ecology of automaticity: Toward establishing the conditions needed to produce automatic processing effects. *American Journal of Psychology*, 105(2), 181–199.
- Bayliss, A.P., Bartlett, J., Naughtin, C.K., & Kritikos, A. (2011). A direct link between gaze perception and social attention. *Journal of Experimental Psychology: Human Perception and Performance*, 37(3), 634–644.

¹³ Thank you to an anonymous referee for helpful feedback on an earlier version of this paper.

- Beck, J. (2018). Marking the perception-cognition boundary: The criterion of stimulus-dependence. *Australasian Journal of Philosophy*, 96(2), 319–334.
- Bigelow, J. & Pargetter, R. (1987). Functions. *Journal of Philosophy*, 84, 181–197.
- Buckner, C. (2014). The semantic problem(s) with research on animal mind-reading. *Mind & Language*, 29(5), 566–589.
- Burra, N., Hervais-Adelman, A., Kerzel, D., Tamietto, M., de Gelder, B., & Pegna, A. J. (2013). Amygdala activation for eye contact despite complete cortical blindness. *The Journal of Neuroscience*, 33(25), 10483–10489.
- Butterfill, S. A. & Apperly, I. A. (2013). How to construct a minimal theory of mind. *Mind & Language*, 28(5), 606–637.
- Calder, A.J., Beaver, J., Winston, J. S., Dolan, R. J., Jenkins, R., Eger, E., & Henson, R. N. A. (2007). Separate coding of different gaze directions in the superior temporal sulcus and inferior parietal lobule. *Current Biology*, 17, 20–25.
- Calder, A.J., Jenkins, R., Cassel, A., & Clifford, C. W. G. (2008). Visual representation of eye gaze is coded by a nonopponent multichannel system. *Journal of Experimental Psychology: General*, 137, 244–261.
- Capozzi, F., Cavallo, A., Furlanetto, T., & Becchio, C. (2014). Altercentric intrusions from multiple perspectives: Beyond dyads. *PloS ONE*, 9(12), e114210.
- Cole, G, Atkinson, M. A., D'Souza, A. D. C., & Smith, D.T. (2017). Spontaneous perspective taking in humans? *Vision*, 1(17): doi:10.3390/vision1020017.
- Cole, G. G., Atkinson, M., Le, A. T. D., & Smith, D. T. (2016). Do humans spontaneously take the perspective of others? *Acta Psychologica*, 164, 165–168.
- Cole, G. G., Smith, D. T., & Atkinson, M. A. (2015). Mental state attribution and the gaze cueing effect. *Attention, Perception & Psychophysics*, 77(4), 1105–1115.
- Conway, J., Lee, D., Ojaghi, M., Catmur, C., & Bird, G. (2017). Submentalizing or mentalizing in a level 1 perspective-taking task: A cloak and goggles test. *Journal of Experimental Psychology: Human Perception and Performance*, 43(3), 454–465.
- Cummins, R. (1975). Functional Analysis. *Journal of Philosophy*, 72, 741–765.
- Dalmaso, M., Castelli, L., & Galfano, G. (2017). Attention holding elicited by direct-gaze faces is reflected in saccadic peak velocity. *Experimental Brain Research*, 235(11), 3319–3332.
- De Souza, W.C., Eifuku, S., Tamura, R., Nishijo, H., & Ono, T. (2005). Differential characteristics of face neuron responses within the anterior superior temporal sulcus of macaques. *Journal of Neurophysiology*, 94, 1252–1266.
- Driver, J., Davis, G., Ricciardelli, P., Kidd, P., Maxwell, E., & Baron-Cohen, S. (1999). Gaze perception triggers reflexive visuospatial orienting. *Visual Cognition*, 6(5), 509–540.
- Ewbank, M. P., Jennings, C., & Calder, A. (2009). Why are you angry with me? Facial expressions of threat influence perception of gaze direction. *Journal of Vision*, 9(12), 1–7.
- Firestone, C., & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for “top down” effects. *Behavioral and Brain Sciences*: doi: 10.1017/S0140525X15000965
- Flavell, J. H. (1977). The development of knowledge about visual perception. *The Nebraska Symposium on Motivation* (Vol. 25, pp. 43–76). Lincoln, NE.: University of Nebraska Press.

- Fodor, J. A. (1983). *The modularity of mind*. Cambridge, MA: MIT Press.
- Freundlieb, M., Sebanz, N., & Kovács, A. M. (2017). Out of your sight, out of my mind: Knowledge about another person's visual access modulates spontaneous visuospatial perspective-taking. *Journal of Experimental Psychology: Human Perception and Performance*, 43(6), 1065–1072.
- Friesen, C. K., & Kingstone, A. (1998). The eyes have it! Reflexive orienting is triggered by nonpredictive gaze. *Psychonomic Bulletin & Review*, 5(3), 490–495.
- Frisby, J. & Stone, J. (2010). *Seeing: The computational approach to biological vision*. Cambridge: MIT Press.
- Furlanetto, T., Becchio, C., Samson, D., & Apperly, I. (2016). Altercentric interference in level 1 visual perspective taking reflects the ascription of mental states, not submentalizing. *Journal of Experimental Psychology: Human Perception and Performance*, 42(2), 158–163.
- Gallagher, S. (2008a). Direct perception in the intersubjective context. *Consciousness and Cognition*, 17(2), 535–543.
- Gallagher, S. (2008b). Inference or interaction: Social cognition without precursors. *Philosophical Explorations*, 11(3), 163–174.
- Gallagher, S., & Varga, S. (2014). Social constraints on the direct perception of emotions and intentions. *Topoi*, 33(1), 185–199.
- Gao, T., McCarthy, G., & Scholl, B. J. (2010). The wolfpack effect: Perception of animacy irresistibly influences interactive behavior. *Psychological Science*, 21, 1845–1853.
- Gao, T., Newman, G. E., & Scholl, B. J. (2009). The psychophysics of chasing: A case study in the perception of animacy. *Cognitive Psychology*, 59, 154–179.
- Gardner, M. R., Hull, Z., Taylor, D., & Edmonds, C. J. (2018). 'Spontaneous' visual perspective-taking mediated by attention orienting that is voluntary and not reflexive. *The Quarterly Journal of Experimental Psychology*, 71(4), 1020–1029. <http://dx.doi.org/10.1080/17470218.2017.1307868>
- Godfrey-Smith, P. (1994). A modern history theory of functions. *Noûs*, 28, 344–362.
- Goldstone, R. L. (1998). Perceptual learning. *Annual Review of Psychology*, 49, 585–612.
- Graham, R., & LaBar, K. (2007). Garner interference reveals dependencies between emotional expression and gaze in face perception. *Emotion*, 7(2), 296–313.
- Hansen, T., Olkkonen, M., Walter, S., & Gegenfurtner, K. R. Memory modulates color appearance. *Nature Neuroscience*, 9(11), 1367–1368.
- Hare, B., Call, J., Agnetta, B., & Tomasello, M. (2000). Chimpanzees know what conspecifics do and do not see. *Animal Behaviour*, 59, 771–785.
- Helton, G. (2018). Visually perceiving the intentions of others. *Philosophical Quarterly*, 68(271), 243–264.
- Herschbach, M. (2015). Direct social perception and dual process theories of mindreading. *Consciousness and Cognition*, 36, 483–497.
- Heyes, C. (2014). Submentalizing: I am not really reading your mind. *Perspectives on Psychological Science*, 9, 131–143.
- Jenkins, R., Beaver, J. D., & Calder, A. J. (2006). I thought you were looking at me: direction specific aftereffects in gaze perception. *Psychological Science*, 17, 506–513.

- Kanwisher, N. (2001). Neural Events and Perceptual Awareness. *Cognition*, 79, 89–113.
- Kuhn, G., & Kingstone, A. (2009). Look away! Eyes and arrows engage oculomotor responses automatically. *Attention, Perception, & Psychophysics*, 71, 314–327.
- LeDoux, J. E. (1996). *The emotional brain*. Simon and Schuster, New-York.
- Lobmaier, J. S., Tiddeman, B., & Perrett, D. I. (2008). Emotional expression modulates perceived gaze direction. *Emotion*, 8(4), 573–577.
- Logan, G. D., & Cowan, W. B. (1984). On the ability to inhibit thought and action: A theory of an act of control. *Psychological Review*, 91, 295–327.
- Mandelbaum, E. (2014). The automatic and the ballistic: Modularity beyond perceptual processes. *Philosophical Psychology*, 28(8), 1147–1156.
- Mandelbaum, E. (2018). Seeing and Conceptualizing: Modularity and the Shallow Contents of Perception. *Philosophy and Phenomenological Research*, 97(2), 267–283.
- Méndez-Bortolo, C., Moratti, S., Toledano, R., Lopez-Sosa, F., Martinez-Alvarez, R., & Mah, Y. H. (2016). A fast pathway for fear in human amygdala. *Nature Neuroscience*, 19, 1041–1049.
- Milders, M., Hietanen, J. K., & Leppänen, J. K. (2011). Detection of emotional faces is modulated by the direction of eye gaze. *Emotion*, 11(6), 1456–1461.
- Milner, A.D., & Goodale, M. A. (1995). *The visual brain in action*. Oxford: Oxford University Press.
- Moors, A., & De Houwer, J. (2006). Automaticity: A theoretical and conceptual analysis. *Psychological Bulletin*, 132(2), 297–326.
- Morgan, E. J., Freeth, M., & Smith, D. T. (2018). Mental state attributions mediate the gaze cueing effect. *Vision*, 2(1), 11: <https://doi.org/10.3390/vision2010011>
- Mumenthaler, C., & Sander, D. (2015). Automatic integration of social information in emotion recognition. *Journal of Experimental Psychology: General*, 144(2), 392–399.
- Nanay, B. (2010). A modal theory of function. *Journal of Philosophy*, 107, 412–431.
- Neely, J. H. (1991). Semantic priming effects in visual word recognition: A selective review of current findings and theories. In D. Besner & G. W. Humphreys (Eds.), *Basic processes in reading* (pp. 264–336). Hillsdale, NJ: Erlbaum.
- Nielsen, M. K., Lance, S., Levy, J. P., & Amanda, H. (2015). Inclined to see it your way: Do altercentric intrusion effects in visual perspective taking reflect an intrinsically social process? *The Quarterly Journal of Experimental Psychology*, 68(10), 1931–1951.
- Nijboer, T. C. W., van Zandvoort, M. J. E., & de Haan, E. H. F. (2006). Seeing red primes tomato: Evidence for comparable priming from colour and colour name primes to semantically related word targets. *Cognitive Processes*, 7, 269–274.
- Nuku, P., & Bekkering, H. (2008). Joint attention: Inferring what others perceive (and don't perceive). *Consciousness and Cognition*, 17(1), 339–349.
- Palmer, S. (1999). *Vision science: photons to phenomenology*. Cambridge, MA: MIT Press.
- Perez-Osorio, J., Müller, H. J., Wiese, E., & Wykowska, A. (2015). Gaze following Is modulated by expectations regarding others' action goals, *PLoS ONE*, 10(11), e0143614. doi:10.1371/journal.

- Perrett, D. I., Hietanen, J. K., Oram, M. W., Benson, P. J., & Rolls, E. T. (1992) Organization and functions of cells responsive to faces in the temporal cortex. *Philosophical Transactions of the Royal Society B: Biological Sciences*, B335 23–30.
- Phillips, B. (2019a). The shifting border between perception and cognition. *Noûs*, 53(2), 316–346.
- Phillips, B. (2019b). The evolution and development of visual perspective taking. *Mind & Language*, 34(2), 183–204.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception. *Behavioral and Brain Sciences*, 22(3), 341–365.
- Qureshi, A. W., Apperly, I. A., & Samson, D. (2010). Executive function is necessary for perspective selection, not Level-1 visual perspective calculation: Evidence from a dual-task study of adults. *Cognition*, 117, 230–236.
- Rhodes, G., Addison, B., Jeffery, L., Ewbank, M., & Calder, A. J. (2012). Facial expressions of threat influence perceived gaze direction in 8-year-olds. *PLoS ONE*, 7(11), e49317: doi:10.1371/journal.pone.0049317
- Ricciardelli, P., Carcagno, S., Vallar, G., & Bricolo, E. (2013). Is gaze following purely reflexive or goal-directed instead? Revisiting the automaticity of orienting attention by gaze cues. *Experimental Brain Research*, 224, 93–106.
- Ristic, J., & Kingstone, A. (2004). Taking control of reflexive social attention. *Cognition*, 94, B55–B65.
- Samson, D., Apperly, I. A., Braithwaite, J. J., Andrews, B. J., & Scott, S. E. B. (2010). Seeing it their way: Evidence for rapid and involuntary computation of what other people see. *Journal of Experimental Psychology-Human Perception and Performance*, 36(5), 1255–1266.
- Santesteban, I., Catmur, C., Hopkins, S. C., Bird, G., & Heyes, C. (2014). Avatars and arrows: Implicit mentalizing or domain-general processing? *Journal of Experimental Psychology: Human Perception and Performance*, 40(3), 929–937.
- Scholl, B. & Gao, P. (2013). Perceiving animacy and intentionality: Visual processing or higher-level judgment? In M. D. Rutherford & V. A. Kuhlmeier (Eds.), *Social perception: Detection and interpretation of animacy, agency, and intention* (pp. 197–229). Cambridge, MA, US: MIT Press.
- Schwitzgebel, E. (2008). The unreliability of naïve introspection. *Philosophical Review*, 117, 245–273.
- Stegmann, U. (2014). Causal control and genetic causation. *Noûs*, 48(3), 450–465.
- Stich, S. (1996). Deconstructing the mind. *New York: Oxford University Press*.
- Surtees, A., Samson, D., & Apperly, I. (2016). Unintentional perspective-taking calculates whether something is seen, but not how it is seen. *Cognition*, 148, 97–105.
- Terrizzi, B. F., & Beier, J. S. (2016). Automatic cueing of covert spatial attention by a novel agent in preschoolers and adults. *Cognitive Development*, 40, 111–119.
- Teufel, C., Alexis, D.M., Clayton, N.S., & Davis, G. (2010). Mental-state attribution drives rapid, reflexive gaze following. *Attention, Perception, & Psychophysics*, 72, 695–705.

- Teufel, C., Alexis, D.M., Todd, H., Lawrance-Owen, A.J., Clayton, N.S., & Davis, G. (2009). Social cognition modulates the sensory coding of observed gaze direction. *Current Biology*, 19(15), 1274–1277.
- Teufel, C., von dem Hagen, E., Plaisted-Grant, K. C., Edmonds, J. J., Ayorinde, J. O., Fletcher, P. C., & Davis, G. (2013). What is social about social perception research? *Frontiers in Integrative Neuroscience*, 6(128), 1–9.
- Tipples, J. (2002). Eye gaze is not unique: Automatic orienting in response to uninformative arrows. *Psychonomic Bulletin & Review*, 9(2), 314–318.
- van Buren, B., Uddenberg, S., & Scholl, B. J. (2016). The automaticity of perceiving animacy: Goal-directed motion in simple shapes influences visuomotor behavior even when task-irrelevant. *Psychonomic Bulletin & Review*, 23(3): 797–802.
- Varga, S. (2018). Toward a perceptual account of mindreading. *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12556
- Wellman, H. (2014). *Making minds: How theory of mind develops*. Oxford: Oxford University Press.
- Westra, E. (2017). Spontaneous mindreading: a problem for the two-systems account. *Synthese*, 194(11), 4559–4581.
- Wiese, E., Wykowska, A., & Müller, J. H. (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), 1–9: <https://doi.org/10.1371/journal.pone.0094529>
- Wilsona, C. J., Soranzob, A., & Bertamini, M. (2017). Attentional interference is modulated by salience not sentience. *Acta Psychologica*, 178, 56–65.
- Wykowska, A., Wiese, E., Prosser, A., & Müller, J. H. (2014). Beliefs about the minds of others influence how we process sensory information, *PLoS ONE*, 9(4), 1–11, e94339: doi: 10.1371/journal.pone.0094339