

Upplevelsen av ögonrörelseintentioner och deras störning i målinriktade handlingar

The phenomenology of eye movement
intentions and their disruption in goal-
directed actions

Maximilian Roszko

handledare / supervisors

Petter Johansson
Lars Hall
Philip Pärnamets

KOGM20

2017-06-09

Masteruppsats i kognitionsvetenskap
Avdelningen för kognitionsvetenskap
Filosofiska institutionen
Lunds universitet

Master's thesis (2 years) in Cognitive Science
Lund University Cognitive Science
Department of Philosophy
Lund University

The phenomenology of eye movement intentions and their disruption in goal-directed actions

Maximilian Roszko

nat11mro@student.lu.se

Many modern psychological theories still assume that humans know about themselves to a wide and accurate extent, in line with the early classical cognitive frameworks. Other competing frameworks, such as dynamic cognition, propose that intelligent behavior can arise from an interaction between the brain, body, and environment, without the need of manipulating explicitly represented knowledge-states. Intentions, i.e. the dispositions to do a specific action, are one such type of mental events that is assumed to be internally monitored to an accurate extent, and is heavily involved in theories modeling goal-directed action. The dynamic framework suggests that there is no reason to assume that humans would naturally have high introspective access to intentions, or are in need of them when making goal-directed actions in the first place. In this study, the extent to which we monitor eye movement intentions, i.e. the intentions to shift one's gaze towards a specific location, and whether they can be expressed in conscious experience, is investigated. A forced-choice decision task was developed where a pair of faces moved systematically across the screen. In some trials, the pair of faces moved additionally as soon as the participants attempted to gaze at the face which was in the front of the movement direction, such that the participants would never see the 'front' face within the center of their gaze. The results of the experiment suggest that humans in general do not monitor their eye movement intentions in a way that allows for mismatches to be consciously experienced and expressed. It was also possible to bias participants into not choosing the alternative that escaped the center of their gaze, if both faces were highly attractive, and doing so without the participants being aware of the manipulation. The results suggest that oculomotor control is another cognitive domain that humans have low access to, and that theoretical models that assume intentions to be central in goal-directed action need to be revised.

1 Introduction

Psychological models attempting to explain how actions are planned and decided often assume to a large extent that people know what their goals, desires, and attitudes are, and subsequently know their intentions prior to the corresponding planned actions (see e.g. Ajzen, 1991; Dickinson & Balleine, 1994). Furthermore, if there was a mismatch between a person's intention and the outcome of the action, the assumption is that she would notice the mismatch, which would diminish her sense of agency (Haggard, 2017; Hommel, 2015). These theories and their assumptions make sense given the classical framework of symbolic cognition, where it is assumed that mental activity, such as problem-solving, is achieved by acquiring, storing, fetching and manipulating knowledge-states (Newell & Simon, 1976; Fodor, 1985; Laird, Newell & Rosenbloom, 1987). Modern theories on the mind regarding goal-directed action that include intentions still make use of these

assumptions that are found in the classical framework. Intentions do make perfect sense in that view, as intentions would be a natural step between a goal-state which is relatively abstract and initiating a specific action which is fully concrete. Relating intentions to actions and agency is also most often implicitly assumed in common sense psychological talk, which increases familiarity to the concept and possibly biases theoreticians into developing models that include intentions. However, modern and more promising views of cognition, such as the dynamic paradigm, do not need to assume that minds perpetually represent knowledge-states explicitly. Instead, minds can be embodied and dynamically coupled with the environment, acting intelligently with the environment directly as a result of the environment interacting with the mind and body (van Gelder, 1995; Wojnowicz et al., 2009; Pärnamets et al., 2015). In these views, it is not clear why actions that appear to be planned necessarily need prior intentions, or intention-like states, that can or can not be brought to awareness, to function (whether an 'intention' refers to a distinguishable neural activity with causal efficacy that is semantically related, or is merely a helpful concept that does not refer to a distinguishable neural activity but still captures some relevance in neural activity, is also an issue, but that will be put aside in this essay; for a relevant discussion see Dennett, 1991). The choice of action can instead result from a complex interplay between the embodied brain, environment and task, where intentions and conscious planning can be invoked in a post hoc explanation to rationalize and explain the decision to act in the given way, whilst not necessarily providing a veridical or helpful causal account of how the decision for that action was generated.

Empirical evidence against the position that we generally know ourselves accurately is also accumulating. It has for example been shown that we can be made to accept false accounts of immediately prior decisions we made ourselves (accepting a decision we did not make as if we made that decision), without becoming aware of those mismatches, and then provide serious explanations for why we made those manipulated choices (Johansson et al., 2005). Johansson et al. demonstrated this by presenting two pictures of faces to participants where they had to decide which face they found more attractive. In some trials, when the participants were given the picture of the face they chose, the experiment leader using sleight of hand gave the participants a picture of the opposite face. Most participants did not become aware of this mismatch, and then proceeded to confabulate as to why they chose the face they held in their hand as if that was the face they actually preferred.

It has also been reviewed and found that we are often unaware of the stimulus that influenced a response from us, and we are often unaware of the response itself, and we are also unaware that the stimulus caused a specific response from us (Nisbett & Wilson, 1977). Based on such evidence, and the fact that we often misattribute causes for other people's ac-

tions, it has been argued that we self-interpret our own mental states in the same fashion as we interpret other people (Carruthers, 2011). As Carruthers has argued, the ability to introspect might have primarily evolved as we learned to interpret other people's mental states, with our ability to interpret mental states being approximative at best. If that were the case, models assuming that humans know themselves well, and make use of information about themselves to make decisions (such as goal-directed action models that centrally include explicit intentions, attitudes or beliefs), would need to be revised in light of the extent to which humans actually know themselves.

Nevertheless, it evidently is the case that we can become aware of our goals, desires and intentions, and we can plan and then execute such plans consciously if we want to. We can compare an intention to lift our arms, and the result of our arms raising, and was it the case that our arms would not move according to how we imagined them to, it would strike our sense of agency and cause doubt in our perception of control (Haggard, 2017). This dichotomy between our ability to bring things such as intentions and attitudes into awareness and using these to plan actions, whilst also often exhibiting false reports of what caused our actions, and not necessarily being in need of intentions to produce appropriate planned behaviors in the first place, suggest that there is a need to explore the role of intentions and possibly revise models that require them.

Do we monitor our eye movement intentions?

A specific domain which has received very little research regarding the goal-directed action models are intentions of the oculomotor kind, which is surprising as eye movements are highly implicated in decision-making, and should therefore be taken into account when modeling goal-directed action. Eye movements have been shown to reflect ongoing thought processes and behavioral goals (Yarbus, 1967), and to be tightly coupled with visual attention (Deubel & Schneider, 1996; Carrasco, 2011); and manipulating people's gaze behavior in real-time can affect preference formations for stimuli and dynamically alter the decision process (Pärnamets et al., 2015; Shimojo et al., 2003; Krajbich, Armel & Rangel, 2010). On top of that, the oculomotor system has been extensively researched and modeled, so that much is known about how and which brain areas are involved in controlling eye movements (Girard & Berthoz, 2005; Sparks, 2002). Thus, as eye movements can be measured accurately, they serve as good models for studying goal-directed movements (Sparks, 2002).

Following goal-directed action models, one would assume that for every top-down planned eye movement, there is a corresponding intention that could be brought to awareness should a person want to. As far as I am aware, no study has investigated whether this is the case yet. Typical visual experience is highly stable, with eye movement control seemingly being operated smoothly and automatically, suggesting that even 'planned' eye movements are often processed automatically outside conscious thought. Yet it certainly is the case that one can intend to shift one's gaze towards a target within one's visual field if one wants to, so it seems that oculomotor intentions are possible mental events.

The instantiating of oculomotor intentions that would be relevant in goal-directed action models should occur for example when a person is tasked with moving her gaze to a specific location or object that she is aware of. It is then possible to

test if oculomotor intentions precede goal-directed eye movements, by tasking participants in ways that would require them to make such planned eye movements. To investigate whether such oculomotor intentions are monitored, one could manipulate the outcome of the participants' eye movements, such that the object they were supposed to move their gaze towards moved away during the eye movement. The classical cognitive views mentioned above then suggest that a person exposed to such conditions would become aware of the mismatch between their intention to move their gaze towards a specific location or object, and the result of their eye movement, which did not succeed in reaching the target. One would only need to make sure to mask the manipulation of the object movement, such that it does not become obvious that a manipulation is occurring, as that would confuse any interpretation of the subjective experiences of the participants. The trick is that during saccades, visual perception is limited such that it is possible to mask movements of objects (Beeler, 1967; Bridgeman, Hendry & Stark, 1975), although if the target simply moved once as the participant shifted her gaze towards it, the participant could become aware of the change as a result of the target not being in the same location relative to the reference frame, not because she was monitoring her eye movement intentions and the potential error signal that occurred. Therefore, an experimental setup needs to circumvent any such effect such that mismatches between the proposed intentions and the outcome can be enhanced in the participants.

This project explores that which takes effect when such mismatches between oculomotor intentions and the result of the eye movements occur. This was done by developing a decision task which sometimes is secretly gaze-contingent. By making participants' eye movements fail repeatedly in some trials (whilst the participants believe that the target they are trying to view is moving for other reasons), it follows from goal-directed action models that they would likely become aware of their incapacity to move their eyes correctly towards a target, even if they knew that they would have difficulties viewing the target for other reasons. If participants do not become aware of oculomotor failures relative to the oculomotor intentions, it might be the case that such intentions do not necessarily operate normally within the oculomotor system, even when goal-directed eye movements are required. Or, it is possible that the intentions cannot be brought to awareness unless explicitly being made to think about them. Either way, were it the case that the failure of oculomotor intentions would not be brought to awareness, it would imply that goal-directed action models would need to be revised. It would also mean that there exists an additional way of manipulating people without them becoming aware of it, if such oculomotor intentions normally would not be monitored. If that were the case, new ways of investigating the human mind would open up. In a more dystopian outlook, it could also mean that billboards with eye trackers incorporated could force naive people to view certain products by following their gaze, biasing them into buying the products. On the other hand, if participants would notice errors in the expected outcome of their eye movements, it would strengthen the view that people prepare themselves for outcomes prior to executing the actions, even for eye movement actions.

Goal-directed models of action

A representative goal-directed action model can be seen in Fig. 1, adapted from Ajzen's theory of planned behavior (1991)

(with regards to these kinds of models, goals are desire-states of how we want the world to become through our actions, that can be brought to consciousness, and vary in complexity and temporality; goal-directed actions are actions that are guided by these desire-states). Within this model, the planning and deciding of actions are directly influenced by one's *intentions* (not all models of this kind make intentions explicitly involved), whose strength is thought to predict the likelihood of engaging in its corresponding action (Ajzen, 1991). The type of intentions being focused on is that which is followed by an action within a short window i.e. a few seconds, not long-term intentions such as intending to watch a documentary in the evening. Such short-term intentions are furthermore influenced by a range of *evaluations*. Evaluations are influencing factors such as attitudes towards doing a specific behavior, norms about what is appropriate to do, and importantly, perception of control, which ties in with our sense of agency. These evaluations in turn are influenced by various related belief states.

Intentions are central in this theory of reasoned action, as they are thought to capture the motivational factors that guide behavior (Ajzen, 1991). But it is important to point out that there are also non-motivational factors that affect whether a given intention is expressed in behavior, such as resources (e.g. energy, time, money) and opportunity, and these are also taken into consideration by the model. Thus, both motivational and non-motivational factors interact with one's perceived behavioral control to determine whether a given behavior will be expressed.

These kinds of models are highly similar to models describing how the sense of agency is acquired in humans, which makes them targets of scrutiny as well. A prominent model regarding how the sense of agency is manifested does similarly to Ajzen's theory of planned behavior assume that an action is preceded by an intention, which causes a motor plan to take shape (Haggard, 2017). Importantly, as the plan is processed, a 'forward model', which is a prediction of how the expression of the plan will affect the world and subsequently our perception of the world, is analyzed by a 'comparator' module, together with the subsequent perception of the world, as the action has expressed itself in the environment. If there is a low discrepancy between the forward model and the outcome, high sense of agency is maintained, while the greater the mismatch, the less agency one feels (Hommel, 2015). This model would predict, in an experiment where one's actions do not lead to the predicted outcome, that one's sense of agency would diminish, including for goal-directed oculomotor actions.

The neurophysiology of goal-directed actions and intentions

Examining the neurological basis for goal-directed actions and intentions, the brain areas that are typically highlighted in goal-directed action are the basal ganglia (BG) and the pre-frontal cortex (PFC) (Buschman & Miller, 2014; Verschure, Pennartz & Pezzulo, 2014; Lisman, 2014). The BG are implicated in both habitual and goal-direction action control (Redgrave et al., 2010), and seen as enabling actions that are desired, and inhibiting those actions that are not desired and in competition with the desired (Prescott, 2008). The PFC on the other hand is involved in many higher-order cognitive functions, those relevant for goal-directed behavior are for instance representing task- and action-spaces (Verschure, Pennartz & Pezzulo, 2014), and supporting the BG in making appropri-

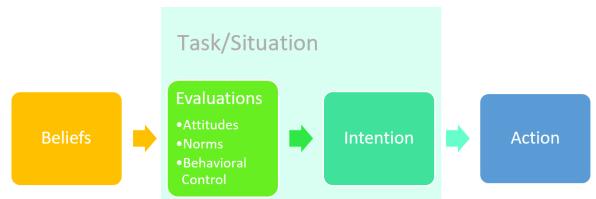


Figure 1. The theory of reasoned action by Ajzen (1991). The choice of actions is thought to be directly influenced by one's intentions which capture the motivational factors that influence a behavior. Intentions are in turn influenced by evaluative factors such as attitudes, perceived norms and perceived behavioral control. The evaluative factors in turn are influenced by a range of beliefs as beliefs about one's attitudes and beliefs about one's control of body and world.

ate task-selections given the environmental and social context. This is achieved through the high interconnectivity between the BG and PFC (Prescott, 2008).

Regarding intentions, it has been found that changes in electrical potential prior to executing planned actions can be ascribed to areas in the frontal and parietal cortex (PC), areas that are involved in motor control (Gilron, Simon & Mukamel, 2015, pp. 105-6). Since intentions come before action execution, it makes sense that activity prior to action execution is designated as intentional, although it is difficult to distinguish such neuronal activity from activity that corresponds to action planning. The supplementary motor area (SMA), pre-SMA, and the premotor cortex, are included in the parts of the frontal and parietal cortex that have been highlighted as important areas for intentions, and the SMA and premotor cortex are high contributors for motor control themselves. Neuronal activity with regards to planned intentional actions has also been found in the parietal reach region (PRR, located in the posterior PC) in macaque monkeys, prior to action execution (Scherberger, Jarvis & Andersen, 2005). In this case the macaques were trained to either saccade or reach towards a point, but not to execute the action until a cue is presented. The activity measured was the changed and maintained activity prior to and during the wait-period when they were not allowed execute the action. That unique activity could then be ascribed as intentional. Interestingly, they found much weaker activity in these areas (PRR, posterior PC) when the planned actions were saccades as compared to arm reaches, presenting a possibility that the conscious saliency of intentions for humans might differ between the different types of intentions one can have, if the same difference in activity is found in the human brain.

Additionally, intentional disorders such as anarchic hand syndrome, where affected patients sometimes perform goal-directed actions whilst denying the intention to do so, are often associated with disturbances with the SMA (Della Sala & Marchetti, 2005). This suggests that the SMA is heavily involved with the feeling of intentions.

The oculomotor circuitry

The oculomotor system does also make use of the BG and PFC in controlling non-reflexive eye movements, suggesting that goal-directed eye movements are not substantially different with regards to brain processing as compared to other kinds

of motor control. Here it is also thought that the BG are involved in selecting actions, and it does so through the help of the PFC (Girand & Berthoz, 2005). For example, the dorsolateral PFC supports a spatial working memory by storing target positions, and is involved in motor planning and selecting responses in relation to goals (Lappi, 2016). The frontal and supplemental eye field receives relevant input from the dorsolateral PFC and parietal eye field via the posterior parietal cortex (which is involved in visuospatial attention), which then engages in selection of action by interaction with the BG and commanding the superior colliculus (SC) (Girand & Berthoz, 2005; Lappi, 2016). Shifting the gaze with rapid eye movements (saccades) is then directly controlled by the saccadic burst generators (which are located in the SC) via the brain stem nuclei, whilst the cerebellum provides with additional regulatory activity (Girand & Berthoz, 2005; Lappi, 2016).

One of the primary roles of the oculomotor system is to gather visual information effectively. Given how the eyes are built, with high visual acuity being restricted to a small part of the retina (the fovea), forces us to rapidly rotate the eyes every time we need detailed visual information from another visible location. This is demonstrated as the typical gaze pattern humans exhibit in everyday situations, where short fixations are interspersed with rapid eye movements (saccades). A challenge with this setup is that rapid eye movements cause the visual information to smear on the retina, such that we effectively become blind during these short eye movements (Deubel, 2004). Another problem is that the light reflection from the world hits the retina differently each time the eye is displaced, even if the objects which reflect the light are stationary. Unless the brain dealt with these neurocomputational problems, we would perceive smears each time we moved our eyes, and the world would look as if it jumped each time. Typical subjective visual experiences tell a different story, one of the world as stable and continuous, so it evidently is the case that the brain processes the rapidly shifting visual information effectively, otherwise we would have clearly noticed it. It should be noticed though that the visual processes making sure the visual world is stable are unconscious, so it is possible that eye movement control is itself largely controlled unconsciously even when eye movements directed by conscious goals are made.

How exactly the brain solves the equations leading to the rise of conscious visual percepts that are stable, in the midst of rapid changes in visual information, is currently under debate. One mechanism thought to provide with stability is the use of corollary discharges (CD), which are motor command copies that are thought to be sent to the visual systems in preparation of the rapid shifts soon to arise as the motor system engages in shifting the eyes (Sommer & Wurtz, 2008). Another factor is the saccadic suppression taking place, which is the dampening of processing new visual information during an interval shortly before, during and after a saccade, that hides the visual disruption (Kowler, 2011). Yet another factor is the occurrence of certain visual neurons that remap their receptive fields, to the area of the retina that currently processes the visual information that is predicted to be processed by the remapping neurons after the saccade, which prepares the neurons for the oncoming change in visual information (Melcher & Colby, 2008). It has also been argued, in an evolutionary fashion, that the visual system assumes the world to be stable during eye movements unless there is evidence of the contrary, so that processing new visual information is facilitated by simply maintaining the pre-

vious processing to a certain degree (O'Regan, 1992).

Together, these mechanisms are thought to help create a stable visual percept of the world, even though most people are unaware of that these processes are taking place and are necessary for a stable visual world. Such cognitive blindness indicates that our metacognitive monitoring of the oculomotor system is limited in some degree, or mostly inactive under normal circumstances. Evidence supports this claim as it has been found that we are not better than chance at identifying our own eye movements from scan paths, as compared to someone else's eye movements in the same task (Clarke et al., 2016). However, it is evident that if one would want to, one could attend one's eye movements and the decisions that control upcoming gaze shifts. It is therefore relevant to explore to what extent internal oculomotor signals are monitored consciously.

Eye movements and their role in decision-making

It is generally assumed that one's visual attention is bound with one's gaze, as it is not possible to shift one's gaze without shifting one's visual attention, although it is possible to shift one's attention without shifting one's gaze (Deubel & Schneider, 1996; Carrasco, 2011). This is partly the reason for why the gaze behavior tells much about the thought processes of a person, and why it is worth investigating both on-line as the eye movements are occurring and off-line from recordings. It also seems to be the case that eye movements are highly purposeful and task dependent (goal-directed but not necessarily consciously monitored) in natural activities, where saccades are made towards targets at 'just the right time' when one needs to engage with them (Triesch et al., 2003), which is another reason for studying eye movements as a window into the mind. However, caution needs to be taken when interpreting eye movement experiments recorded in laboratory settings, if the setup is highly artificial with participants' heads having to be restrained. In naturalistic environments, the head and body are most often not still, even when fixating on something with the eyes. In such instances, compensatory eye movements are engaged to keep the gaze fixed at the target. The result is that fixations in naturalistic settings are not identical to fixations in the lab, as in the lab the head itself is fixated to make it easier to measure the eyes (Lappi, 2016).

Nevertheless, eye movements play a tremendous role in providing visual information even in artificial settings, even to such a degree that eye movements can be manipulated in several ways to alter decision-making in real time. One way to bias decision-making is through increasing exposure towards specific stimuli, for instance by limiting the amount of time a person gets to view a stimulus, while decreasing exposure to others (Montoya et al., 2017); although exactly how the exposure effect alters subjective values is currently debated as there are multiple models attempting to explain it. In general, it is thought that as one is being exposed to a stimulus, more information is collected about it, which dynamically feeds one's preference or disposition towards that stimulus. This effect can be seen to both increase one's liking toward a target, if the target is appealing in the first place, or decrease preference if the target is aversive (Armel, Beaumel & Rangel, 2008).

Preferences have been biased in experiments using eye gaze dynamics for faces (Shimojo et al., 2003), moral positions (Pärnamets et al., 2015), and abstract shapes (Nittono & Wada, 2009), among other things. Shimojo and researchers did it by manipulating the amount of orientations towards certain alter-

natives, which increased the exposure towards those alternatives, and according to the researchers the process of orienting towards a target stands itself as a factor that biases preferences. Pärnamets et al. influenced preferences by dynamically terminating the viewing process according to differences in accumulated viewing time between two moral positions, and Nittono and Wada did it purely by controlling stimuli exposure time, increasing it for randomly selected stimuli.

This empirical evidence suggests that decisions are dynamically processed in real-time, whereby available information and environmental coupling is continually updating and influencing the decision process, and vision seems to be a very influential channel for feeding available information.

Predictive processing

Another framework that strikes the middle-ground between viewing cognition as dynamic, but positions predictive states (which intentions arguably could be classed as) as central in perception and action, is the *predictive processing* framework (PP). Instead of viewing the brain as a system which constructs reality and acts on it according to information gathered through the senses, the PP framework assumes the brain to do the opposite: it predicts top-down how the world should be, given the experience and knowledge of the perceiver, and the situated environment the perceiver is located in (Clark, 2015). This assumption turns out to explain many questions pertaining to how it is possible for the brain to process the vast amounts of information the senses provide each moment, and how the brain manages to interact with the world through the body seemingly without complication. Such a way of operating on perceptual information allows the brain “to select frugal, action-based routines that reduce the demands on neural processing and deliver fast, fluent forms of adaptive success” (Clark, 2015, p. 1).

The general principle by how this process operates is that the brain perpetually constructs what is called ‘generative’ or ‘forward’ models (Friston, 2012). For instance, the brain networks responsible for perception, are viewed as hierarchical networks, where the topmost layers provide abstract predictions (forward models) for the level below, and the further down you go in the hierarchy, the more concrete both spatially and temporally the predictions become (Clark, 2015, pp. 29-31). A level that receives predictions provides with error signals back up based on how well the predictions it receives matches the error signals that layer itself received from the predictions it sent to a layer below. The middle layers essentially compute what the predictions it received failed to include or exclude, from information that it received from another level further down the hierarchy. This process repeats itself all the way down to the sensory neurons, who are not thought to provide signals solely based on the stimulation they receives from the world, but also on the predictions they received from their corresponding level above (see Lupyan & Clark, 2015, for a more thorough explanation together with figures).

Such a way of processing sensory information has been shown for the visual system for instance, where it has been found that ganglion cells do not signal the raw visual information they get, but compute a difference between the raw information and the predicted information they should receive (Hosoya et al., 2005). Regarding higher-level functions, this framework also excels at explaining neurological diseases like autism and schizophrenia, among other things. Countless stud-

ies have shown how patients with schizophrenia have a reduced ‘mismatch negativity signal’, which has been interpreted as a weakened error signal (Todd et al., 2012). If a predictive system lacks the ability to functionally check the predictions and provide with error signals when the bottom-up signals provide with information that does not match with the predictions, the predictions stand unregulated, which explains how schizophrenic patient can hallucinate objects and sounds even when no such signals originate from the lower sensory layers. It also explains how it is that we cannot tickle ourselves: The tactile stimulation induced by self-stimulation is down-regulated as a result of the forward model of the motor system, which not incidentally is something schizophrenic patients are less capable of doing. And as the theory explains, not being able to properly regulate forward models and sensory signals is an explanation for why schizophrenic patient have enhanced capabilities for self-induced tickling (Blackmore, Wolpert & Frith, 2000).

Crucially, predictive processing applies to action as well. It is through action that we bring about the changes in sensory stimuli which need to be predicted in the first place (Clark, 2015, pp. 6-7). Our motor plans allow us to form better predictions of the sensory outcome in the near future, which again ties in with the ability to experience a sense of agency, since when we cannot predict the outcome of our actions, there is a loss in the sense of being able to do things according to how one wants them to go, which is part of what being an agent is. Similarly to the corollary discharges (motor command copies) that were mentioned for the oculomotor system, the same is true for the motor system in general. Not surprisingly, schizophrenic patient exhibit a weakened sense of agency, which is linked to disturbances in their CDs; moreover, schizophrenic patients exhibit multiple types of oculomotor disturbances, that are thought to be a result of the CD disturbances (Thakkar, Diwadkar & Rolfs, 2017).

The PP framework predicts that there should be some sort of error signaling present when oculomotor predictions fail, such as for instance when *intending* to make a specific eye movement with a result in mind, but achieving the wrong results. This has exactly been shown, where saccade-contingent changes elicited error signals, that were found in the visual cortex using electroencephalography (Ehinger, König & Os-sandón, 2015). The question though is whether these error signals can be perceived consciously or not, as that has not been demonstrated for oculomotor error signals as far as I am aware. It is also not clear how error signals relate to intentions, and if different error signals affect awareness differently. It might be that when subject to severe circumstances with perpetual mispredictions, causing repeating instances of error signals, one becomes more vigilant, with other conscious changes relating to that, which might cause one to notice intentional mismatches eventually. However, it is not given that humans are able to directly experience such error signals. Noticing mismatches between intentions and outcome might be done through indirect means, such as through reasoning with other sets of consciously available information.

Hypotheses

Although the project is exploratory in nature, several hypotheses can be tested given the available theory. The goal-directed action models postulate a specific role for intentions, in that they are a required step before planned actions can execute.

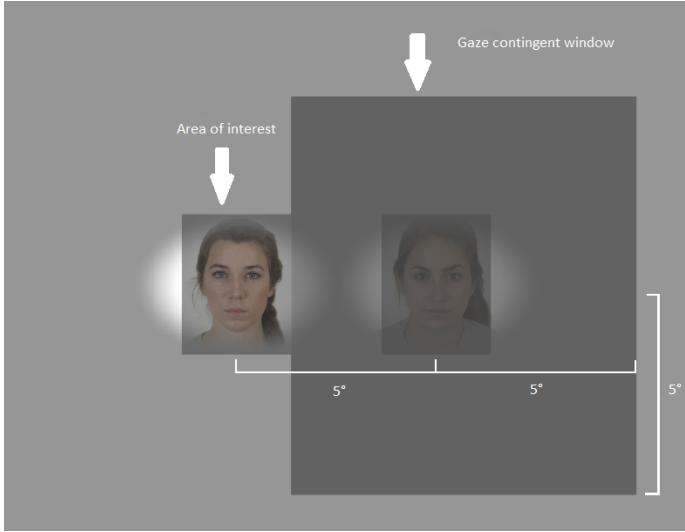


Figure 2. The presentation of a pair of faces in a decision trial (only manipulated trials included gaze-contingent windows). If the movement of direction was to the right, the gaze-contingent window was centered on the face to the right, which was then called the 'front' face, and vice versa if the movement of the pair was to the left. The faces were separated by 5 degrees, although when transitioning from the 'back' to the 'front' face, a movement would be triggered when leaving the area of interest which was centered on the 'back' face, the barrier being approximately 3.6 degrees from the center of the 'front' face. The grayed areas on this figure were not visible for the participants.

Thus, if participants are being made to produce goal-directed eye movements towards a target (goals that are consciously available, e.g. "I want to look at that face"), but continually fail in reaching it, the theory suggests that participants would notice this mismatch between intention and results, and would readily express it spontaneously in an interview after. It should be noted however, that the participants need to expect that the stimuli will move as part of the task, but not that the stimuli will move as the participants attempt to look at them. This allows for the participants to notice the manipulation only if they monitor their eye movements intentions.

Another hypothesis is that according to the models on how the dynamics of gaze can bias decisions, participants will be less likely to choose the randomly blocked faces, as we move those faces away from the participants' gaze every time they try to view those faces. This predicted bias is assumed to partly be due to choosing the alternative one is more exposed to. And since there is greater uncertainty about the blocked face as one has not seen it properly, it is likely that one will not choose it because of that. However, if the faces are unappealing, it is also predicted that the participants will be less likely to choose the face they are forced to view, even when the other face keeps moving away from their gaze, as increased exposure to unappealing stimuli will increase confidence in not liking that alternative.

It is also assumed that the participants will have worse visual processing of faces that are made difficult to perceive through the gaze-contingent manipulation, those faces which participants will essentially be blocked in viewing. But to what extent the perception of the blocked faces will be is unknown, as in this experiment the participants will be given free control over the amount of time spent on each decision trial. By allowing participants to freely control the exposure time, even

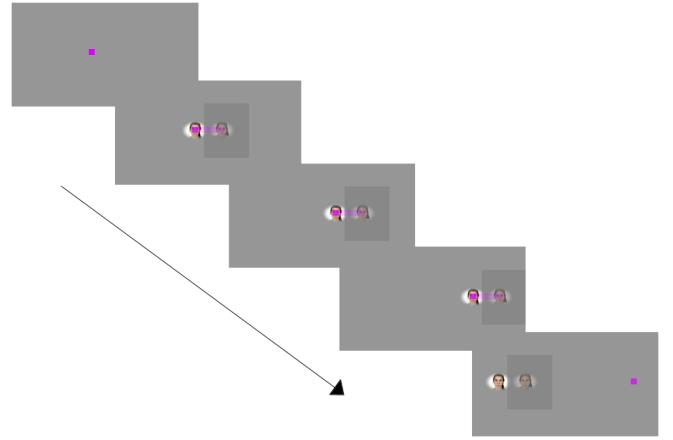


Figure 3. The beginning of a typical manipulated decision trial, with a purple dot representing a schematic gaze position and movement. A fixation marker could appear at any of 8 evenly spaced horizontal locations, which would draw the participants gaze towards it. After a random time between 1 and 2 seconds, the pair of faces would appear, with the 'back' face always appearing where the fixation marker (and most likely the participant's gaze) was located. The 'front' face appeared either to the left or right of the 'back' face depending on the trial's movement direction, which was randomly determined. The faces together jumped 5 degrees in their movement direction after a random amount of time between 42-292 ms, or when the participant had triggered the gaze contingency by moving their gaze within the gaze-contingent window in manipulated trials. When the pair of faces reached the end of the monitor, such that another jump would locate one of the faces 'outside' of the screen, they instead reappeared on the other end of the monitor and began anew in the same movement direction. The grayed zone around the 'front' face was not visible to the participants.

when not being able to view one of the faces in a trial, the participants effectively signal when they are satisfied with making the decision each time they decide which face they find more attractive, serving as an interesting factor when comparing with their actual ability to discern whether they recognize the face in a memory task later. This will be explored by a memory test after the decisions have been made, where it is expected that participants' recognition of faces in a set of previously shown and not shown faces will be significantly reduced for the faces that were blocked from view.

2 Method

The purpose of the experiment was to explore to what extent oculomotor intentions are monitored, and how that could express itself phenomenologically. Therefore, a forced decision task was constructed where the alternatives in the decision, in this case faces, jumped repeatedly across the screen, in order to have a subset of these trials where participants would automatically fail in reaching one of the faces with their gaze, as that alternative would always move away as they made a saccade towards it. The movement of the faces in normal trials was designed so as to effectively mask trials that were gaze-contingent, by resembling the pattern of movement of the stimuli during gaze-contingent trials. Hence, the stimuli in manipulated trials would move in the same random way as in trials without manipulation if the participants had their

Table 1. The set of interview questions, translated into English from Swedish. The primary questions were always asked, while the follow-up questions interjected whenever a participant expressed difficulties in seeing the faces during the experiment. Question 5 in the primary set was not asked if previously answered in the follow-up questions.

Primary questions	Follow-up questions
<ol style="list-style-type: none"> 1. How do you feel about the experiment? 2. Did you think about anything in particular during the experiment? 3. How did it feel to make the decisions? 4. Was there anything in particular that you thought about regarding the decisions? 5. Did you evolve any strategies for decision-making? 6. How do you feel about the faces? 7. Was there anything in particular that you thought about the faces? 8. Did you notice anything weird about the experiment? 9. If you had to guess about an undisclosed purpose of the experiment, what would you guess? 	<ol style="list-style-type: none"> 1. Why do you think it was like that? 2. Was it difficult to see the faces in a particular position? 3. Could you fixate both faces? 4. Could you see both faces? If not, what strategy did you use when making the decisions? 5. Did the decisions feel free or dejected when it was difficult to see?

eyes closed, but if the participants tried to view the designated 'blocked' face, they would trigger the gaze contingency, which would cause the faces to jump before the random wait-period had expired.

An interview after the experiment peered into the participants' phenomenology, in an attempt to explore the degree to which the repeated failures of oculomotor intentions are self-monitored. The presumption was that if the participants were monitoring their oculomotor intentions, they would notice that their eye movements caused the stimuli to move and that they were unable to view one of the alternatives no matter how hard they tried.

A memory phase followed the decision phase, asking them whether they had seen the currently presented face, and if so, whether they think they preferred that face in the decision it was part of. By doing so, the participants' recognition and source memory for blocked and non-blocked faces could be compared. The memory test would also serve as a control that the manipulation worked.

Additionally, with this setup, we would also investigate how decisions are affected when only one alternative is clearly visible out of a pair, even though the participants would be granted unlimited time to make their decisions.

Design

The experiment was a within-subject design, and consisted of 60 decision trials and 180 memory trials. In the decision trials, pairs of faces were presented which rapidly jumped across the screen, and the participants' task was to decide which face out of a pair they found more attractive. The five first decision trials were always without gaze contingency, while the remaining 55 trials could be either gaze-contingent or not with 50% likelihood for either. The participants had unlimited time to make each decision, and after each decision they had to input their degree of confidence for the decision they just made.

The system used to accomplish the manipulation was programmed to block participants in viewing the face that was in the first position of the pair relative to the direction of movement in the manipulated trials. The strength of visual processing in the periphery had thus to be accounted for, as this factor would influence to what degree the participants could 'see' the blocked faces depending on the distance we separated them. A study has shown that it is possible to reliably determine attractiveness between faces at distance of 5 visual degrees (Guo, Lie & Roebuck, 2011), although in those pairs the attractiveness difference was high between the faces. This is one of the

reasons for why the faces in a pair have been matched in attractivity, provided through a norming table supplied by the database used for the faces.

The distance between the faces had also to be taken into account. The visual acuity relative to the fovea drops sharply, such that at a distance of 5 degrees from the fovea, the relative acuity is about 30% (Hans-Werner, 2006). It was thus deemed good enough to separate the faces in a pair at 5 degrees, such that participants could get a sense of the blocked face in manipulated trials, while in fact not being able to see its details better than 30% of acuity compared to the face that was not blocked.

The memory test, in which a single face was presented, asking participants whether they recognize the face, included the 60 pairs of faces that were part of the decision task, with an additional 60 faces that were not previously presented, for a total of 180 faces. Out of the total of 180 available faces, those that were part of the decision phase were randomized for each participant, while the faces always maintained the same pairing. The order of the pairs in the decision phase and the relative position of the faces within a pair (left vs. right) were random as well. The pair of faces could also jump either in a rightward or leftward direction, and the pair could appear on the screen after the fixation marker at any of 7 different positions.

Experimental setup and stimuli

The participants had their heads on a chin rest 80 cm in front of a 27-inch LCD monitor (resolution at 1920 x 1080 pixels) with a refresh rate at 120 Hz. The eyes were measured using the Eyelink 1000 (SR Research, Ontario) that recorded monocularly at 1000 Hz, while the experiment was run on Python 2.7.3, mainly using the PsychoPy module (Pierce, 2007). All eye movement data were recorded online after a nine-point calibration (average measured accuracy = 0.47, $SD = 0.33$).

The faces came from the Chicago Face Database (Ma, Correll & Wittenbrink, 2015). All faces used were frontal-view Caucasian with neutral expression pre-rated for a number of subjective attributes on a 1-7 point scale, and the faces were paired gender-wise and according to their closest proximity in attractivity scores, which were provided by the database. The faces were divided into three attractivity groups, the highest 25% belonging to the 'high' attractivity group, the middle 50% belonged to the 'mid'-group, and the lowest 25% belonged to the 'low'-group. This divide was made to separate the high and the low group as much as possible while maintaining enough

faces in those groups to reliably show an effect for facial attractiveness in the decisions if it exists. The images of the faces were resized to 244 (wide) x 172 (high) pixels, with a raised cosine edge that was applied to provide softness to the images (Fig. 2). The distance between the centers of the images in a pair was approximately 5 visual degrees (230 pixels on the screen at a distance of 80 cm from the eyes, 33 pixels/cm).

The gaze contingency was determined by a few conditions. If the pair of faces were moving in the rightward direction, and if the participant's gaze crossed slightly before the 'left' image border, the stimuli would jump towards the right 5 visual degrees, as that would capture when the participant was making a transition from the 'left' towards the 'right' face. But it could also be triggered if the gaze approached from the other direction and came within 5 degrees of the center of the 'right' image (Fig. 2), as this could happen when the faces had reached the edge of the monitor and spawned at the opposite side, with the participant shifting her gaze across the screen to catch up with the faces. The opposite applied if the direction of the jumps were 'left'. But if the participant did not trigger the gaze-contingent conditions, the movement of the faces continued in the same fashion as normal trials.

The base rate at which the faces jumped in normal trials was dependent on a few conditions designed to mask the fact that the manipulated trials were gaze-contingent. The time a pair would remain in position was a number of frames of equal probability between 20-35, times the length in time for each frame (at 120 Hz each frame was about 8.3 ms), such that the pair could stay in the same position between about 166-292 ms. Additionally, there was a 10% probability that the faces would jump after only 5 frames (42 ms), to produce more jittery behavior that resembles the gaze pattern of participants attempting to view a blocked face in purely gaze contingent conditions.

There was a total of 7 different positions that the pairs could be located on, and as soon as a pair reached the edge of the monitor, they appeared at the opposite side of the monitor and began anew in the same direction as before (Fig. 3).

Participants

31 participants (17 female, 14 male) recruited at Lund University, mostly students (mean age = 26.1 years, $SD = 7.1$), took fully part in the experiment. All participants reported normal or corrected-to-normal vision with contact lenses, as doing the experiment with glasses was not allowed due to recording difficulties. The participants received compensation in the form of a gift voucher valid at the movie theater for participating.

Measures and analysis

To determine whether participants monitored oculomotor intentions consciously to any extent, a set of questions (Table 1) was devised that would probe the participants' subjective experiences regarding oculomotor actions and decision-making, while trying to limit the extent of leading questions that could produce post hoc rationalizations. The set of primary questions was devised to scan for any experiences in the participants that could relate to oculomotor intentions, such as if they felt that their eyes did not allow them to see one of the alternatives, or whether it was impossible to see one of the faces no matter how hard they tried to look at it. If any response from the participants sufficiently seemed as there could be some relevant awareness, a set of follow-up questions interjected wher-

Table 2. The frequency of participants belonging to the each category of awareness.

Degree of awareness	Amount
1. No reflections regarding perceiving the faces clearly	3
2. Experienced difficulties perceiving both faces clearly sometimes	18
3. Suspicious that faces were sometimes manipulated to be faint/blurry/unclear	7
4. Experienced feelings that faces sometimes moved according to eye movements	0
5. Suspicious that faces moved according to eye movements	3

ever that occurred in the primary set of questions, returning back to the primary questions after the follow-up questions were completed.

The degree of awareness was divided into 5 categories specifically related to this experiment: whether the participant made no reflections regarding difficulties in viewing both faces sometimes, whether the participant reflected that it was difficult to perceive both faces sometimes, whether the participant reflected on being suspicious that something was done to the faces making them more difficult to perceive sometimes, whether the participant reflected that she felt as if her eye movements affected the movement of the faces, and whether the participant was suspicious that her eye movements directly affected the movement of the faces. The answers participant provided to the questions asked were then used to categorize the participants accordingly.

To determine how the manipulation and relevant factors affected decision-making and memory of the faces, generalized linear mixed models (GLMM) were calculated using the **lme4** package in R (Bates et al., 2015). Participants' confidence were analyzed by a linear mixed model (LMM) with the same R package. Random effects were modeled as per participant intercepts, and reflecting the fixed effects structure to the closest degree such that convergence was achieved. Significance for fixed effects for the LMM was calculated with the **lmerTest** package.

Signal detection analysis was used to calculate per participant the sensitivity of recognizing faces, both comparing the sensitivity for participants between faces belonging to either decision trials that were manipulated or not, and comparing the sensitivity participants had between faces that were blocked or forced in the manipulated trials. The calculations for d' were done using the empirical probit transform.

The values from the confidence ratings would be used to interpret the other results.

Procedure

Participants were introduced to the experiment, and were told that the purpose of the experiment was to investigate how moving alternatives in a decision affected their decision-making. They were told that several measurements would be recorded, for instance pupil size and reaction time, but nothing about eye movements. They were also told to make their decisions without feeling pressured for time.

After being given the instructions, the participants were monocularly calibrated. After that, the decision trials commenced. Participants had to make 60 decisions, and an ad-

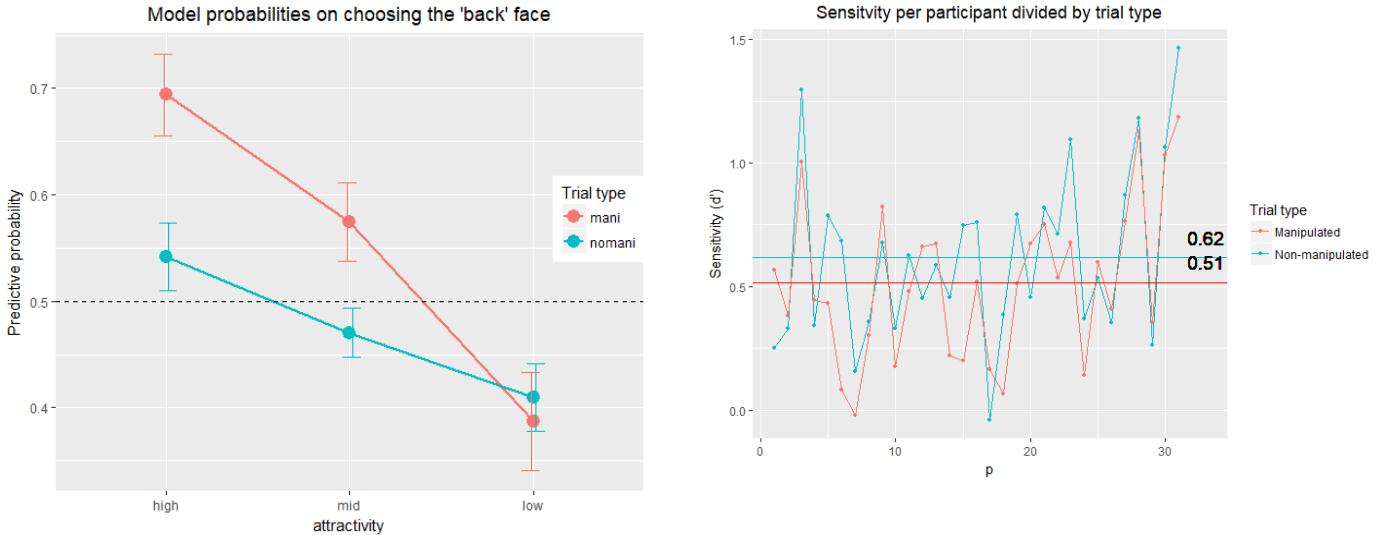


Figure 4. The predictive probabilities and their standard errors of choosing the 'back' face, which in manipulated trials were the 'forced' faces which participants had forced biased viewing times towards. The factors are trial type and attractivity, with their interaction included in the values. There was a significant difference between the trial types and between 'high' and 'low' attractivity, which shows a shift in bias from the 'back' and the 'front' face when the faces were 'low' in attractivity. mani = manipulated trial; nomani = non-manipulated trial.

ditional 60 confidence judgments intervening each decision regarding how confident they felt they were about the decision they just made. In the decision trials, a fixation marker spawned at one of 8 possible positions, and after a random time between 1-2 seconds the pair of faces appeared, such that if the direction was rightwards, the 'left' face spawned at the spot that the fixation marker was at previously. The pair of faces started moving immediately after spawning according to its program, and the participant had unlimited time to make the decision by pressing either 'left' or 'right' depending on the preferred face. After each decision a confidence-scale appeared, where the participant could input their confidence in the previous decision on a continuous scale from 1.00 to 6.00.

After the decision phase, they were instructed about the memory test, which they then proceeded to make 180 trials of. During each trial a face appeared, with instructions asking if they recognized the face, pressing 'left' for yes and 'right' for no. If they answered positively, they were additionally asked whether they think they chose that face in the decision trial it was part of or not. Two-thirds of the faces had been presented previously, but the participants were told that the ratio between old and new faces varied between the participants and would not necessarily be 50/50.

After the two phases, the participants were interviewed. There were 9 primary questions asked, but if a participant expressed some suspicion regarding the movement of the faces, or if they expressed that they could not perceive both faces sometimes, a set of follow-up questions intervened (Table 1).

3 Results

Interviews

The division of participants according to the defined degrees of awareness can be seen in Table 2. Very few participants (3)

Figure 5. The distribution of d' values per participant divided up by trial type. Participants were on average better able to determine whether they recognized faces that were part of non-manipulated trials. The horizontal lines represent the mean for each category.

explicitly noticed the manipulation that took place in the experiment, while most participants (18) did not notice the manipulation but expressed experiencing frequent difficulties in viewing the faces. While there are too few participants in each group to achieve statistically significant differences between the different categories, those participants (3) who made no expressions regarding difficulties in viewing the faces still had observable differences (in line with the whole group) in their confidence, decision time, and number of transitions between the trial type conditions. This indicates that there likely was no technical difficulties which could have reduced the functionality of manipulated trials such that 'blocked' faces were not blocked properly.

The type of response that was frequent for participants who were categorized in group 2 was for instance that "many of the faces blurred together" or that "they moved so fast that you could not see both clearly." What separated class 2 responses with class 3 were that in the latter case participants explicitly stated that they were suspicious to some degree that the faces were manipulated (although not the movement of the faces but their attributes). A typical class 3 response is for instance that "I thought about whether the faces were recycled, and whether they were modified," or that "it felt as if it made up sometimes, as if they were not real people, as if the faces' width were extended." Importantly, participants belonging to class 3 never expressed that their eye movements affected the movement of the faces, although they could have expressed that where they looked first in a trial might have affected which face got manipulated into looking unreal. No participant was categorized to class 4, as those participants who expressed that their eye movements affected the movement of the faces did so with confidence or highly accurate remarks. Those participants who noticed the manipulation clearly expressed for instance that "I tried to understand how it worked, it felt as if when the eyes moved a lot the pictures moved even more," or that "it felt as if the system did that, that you should look at the picture in the back."

Regarding the participants' feelings as decision-makers, 13 participants expressed feeling that they could make free decisions, whilst the rest (18) told that they felt their decisions to be constrained and insecure.

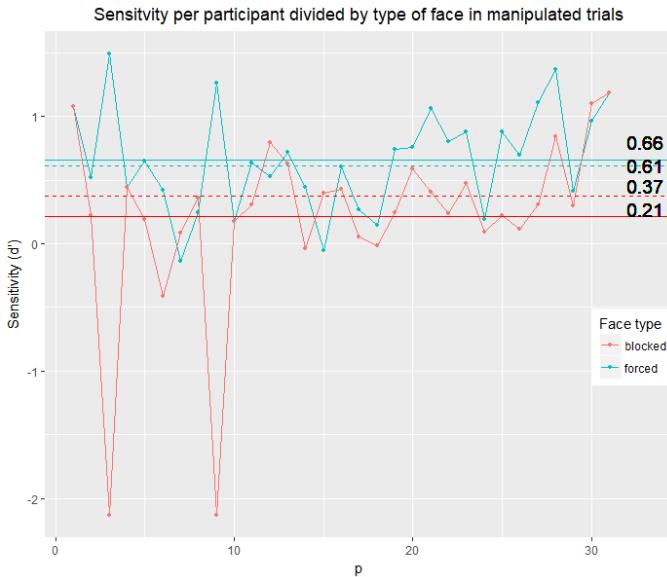


Figure 6. The distribution of d' values per participant divided up by manipulated trials and whether the face was blocked or not in the trial it belonged to. On average, participants were better recognizers of faces that were 'forced' to be viewed, compared to faces that were 'blocked' from being viewed. The filled horizontal lines represent the mean values for each category, while the dashed lines represent the mean values for each category excluding the two outliers.

Quantitative data

The time spent per decision trial differed between the trial types, as the participants spent on average more time on the manipulated trials ($M = 6.41$ s, $SD = 5.12$ s; for non-manipulated trials: $M = 4.69$ s, $SD = 3.65$ s). Participants' average confidence responses were also lower for the manipulated trials ($M = 3.17$, $SD = 1.32$; for non-manipulated trials: $M = 4.01$, $SD = 1.10$). While in the manipulated trials where participants were blocked in fixating the face furthest in the direction movement (the 'front' face), there was also a bias to have spent more time on the face which was last in the direction movement (the 'back' face) for the non-manipulated trials (average time spent on the 'back' face = 1.40 s, $SD = 1.20$ s; average time spent on the 'front' face = 0.77 s, $SD = 0.73$ s; average relative time spent on the 'back' face = 65%, $SD = 13\%$).

Decisions

Whether the participants chose the 'back' or the 'front' face was also modeled with a GLMM. The chosen factors were trial type, attractivity, and their interaction. There was a significant fixed effect for manipulated trials as compared to the non-manipulated trials, $\beta = 0.653$, $SE = 0.216$, $p = 0.00252$ (a positive fixed effect means here that it positively affected choosing the 'back' face), for 'low' attractivity as compared to 'high', $\beta = -0.532$, $SE = 0.183$, $p = 0.00359$, and the interaction between manipulated trial type and 'low' attractivity, $\beta = -0.747$, $SE = 0.283$, $p = 0.00834$. No significant effects were found for 'mid' attractivity, $\beta = -0.285$, $SE = 0.154$, $p = 0.0647$, and the interaction between manipulated trial type and 'mid' attractivity, $\beta = -0.234$, $SE = 0.240$, $p = 0.329$. The model probabilities for choosing the 'back' face can be seen in Fig. 4.

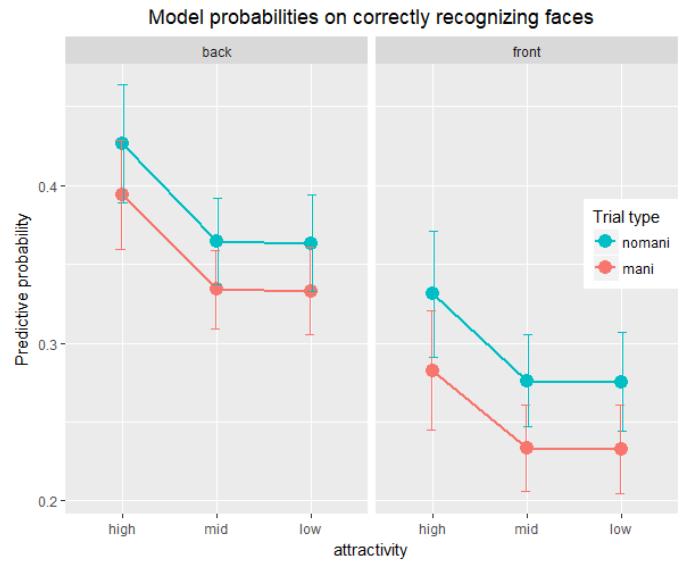


Figure 7. The predictive probabilities and their standard errors of responding correctly to the recognition task depending on where the face was positioned in the decision trial, the attractivity of the face, and whether the face belonged to a manipulated or non-manipulated decision trial. The modeling excluded recognition trials with new faces. mani = manipulated trial; nomani = non-manipulated trial.

Memory

The participants answered correctly whether the faces they saw were part of the decision trials in 2766 out of 5580 trials (49.6%). Out of the memory trials which progressed into the participant deciding whether she thinks she chose the face in the decision it was part of, they were correct in 649 out of 1566 trials (41.4%).

Signal detection analysis revealed that the average recognition sensitivity for faces that were part of manipulated trials was lower ($M = 0.51$, $SD = 0.32$) than the average recognition sensitivity for faces that were part of non-manipulated trials ($M = 0.62$, $SD = 0.35$) (see Fig. 5). Comparing how the recognition sensitivity differed in manipulated trials between face position ('back' position was also called 'forced', 'front' position was also called 'blocked'), it was found that participants were on average less sensitive for blocked faces ($M = 0.37$, $SD = 0.36$) than forced faces ($M = 0.61$, $SD = 0.37$), even when eliminating two extreme outliers that are in favor of a difference (see Fig. 6).

A generalized linear mixed model was calculated to compute the effects of trial type, attractivity, the position of the face and the interaction between the position and trial type, on being able to correctly recognize faces in the memory phase. The significant fixed effects were 'front' position as compared to 'back' position, $\beta = -0.407$, $SE = 0.105$, $p < .001$ (negative fixed effects means here negative effect on correctly recognizing the faces); 'low' attractivity as compared to 'high' attractivity, $\beta = -0.264$, $SE = 0.125$, $p = .0343$; 'mid' attractivity, $\beta = -0.260$, $SE = 0.0958$, $p = .00668$. The fixed effect on trial type which is manipulated as compared to non-manipulated, showed no significance according to this model, $\beta = -0.133$, $SE = 0.103$, $p = .194$, and the interaction between manipulated trial type and 'front' position was not significant either, $\beta = -0.0944$, $SE = 0.154$, $p = .539$. The model probabilities for correctly recognizing the faces can be seen in Fig. 7.

How trial type, attractivity, position and interaction be-

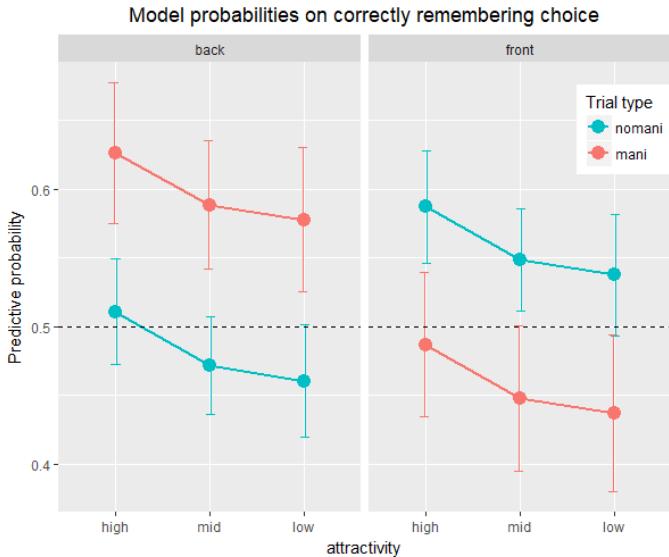


Figure 8. The predictive probabilities and their standard errors of responding correctly to the recognition task depending on where the face was positioned in the decision trial, the attractivity of the face, and whether the face belonged to a manipulated or non-manipulated decision trial. The modeling excluded recognition trials with new faces. mani = manipulated trial; nomani = non-manipulated trial.

tween position and trial type affected the task of answering whether they think they preferred the face in the decision trial it was part of, was also modeled with a GLMM. Significant fixed effects were found for manipulated trials as compared to non-manipulated trials, $\beta = 0.472$, $SE = 0.221$, $p = .0328$ (a positive fixed effect here means that it had a positive effect on answering correctly), and for the interaction between manipulated trials and the 'front' position, $\beta = -0.877$, $SE = 0.390$, $p = 0.0246$. Non-significant effects were found for 'low' attractivity as compared to 'high', $\beta = -0.202$, $SE = 0.169$, $p = 0.233$, 'mid' attractivity, $\beta = -0.157$, $SE = 0.142$, $p = 0.268$, and 'front' position as compared to the 'back' position, $\beta = 0.309$, $SE = 0.212$, $p = 0.146$. The model probabilities for correctly remembering whether preferred the faces they recognized can be seen in Fig. 8.

Confidence ratings

The confidence ratings were modeled according to trial type, attractivity, their interaction, decision time, and number of transitions, with a linear mixed model. Their estimates and standard errors can be seen in Fig. 9. The intercept represents the mean in confidence (from 1 to 6 max) when the trial type was non-manipulated, the attractivity was 'high', and transitions and decision time was 0. Most values are negative except the interactions between trial type and attractivity, and trial type and decision time. It is also the case that these interaction effects were non-significant, for the interaction effect between manipulated trial and 'mid' attractivity, $\beta = -0.109$, $SE = 0.100$, $p = .274$; for the interaction between manipulated trial and 'low' attractivity, $\beta = -0.103$, $SE = 0.118$, $p = .381$; for the interaction between manipulated trial and decision time, $\beta = 0.00259$, $SE = 0.0174$, $p = .884$. There was a significant effect for manipulated trials as compared to non-manipulated trials, $\beta = -0.652$, $SE = 0.158$, $p <.001$, and 'mid' attractivity as compared to 'high' attractivity, $\beta = -0.257$, $SE = 0.0709$,

$p <.001$, and for 'low' attractivity, $\beta = -0.385$, $SE = 0.0836$, $p <.001$, and for decision time, $\beta = -0.0860$, $SE = 0.0208$, $p <.001$, and for transitions, $\beta = -0.0112$, $SE = 0.00398$, $p = .0108$.

4 Discussion

The primary purpose of the experiment was to investigate the extent to which humans monitor their oculomotor control, specifically their intentions to move their gaze towards a target with saccades. Understanding whether it is the case that people generally do monitor the intentions behind their eye movements is interesting because there are conflicting models and frameworks that describe human cognition: The symbolic framework would support the idea that it would be natural for humans to have accurate introspective access to their action-planning, especially when it is goal-directed; while the dynamic side would argue that decision-making and action-control processes, even when goal-directed, are to a large extent dynamic, situated, and outside of conscious control, with introspective abilities being limited in depth and accuracy. Studying the oculomotor system is hence very favorable in trying to shine a light on this conflict, as eye movements are in general automatically controlled, while they can easily be consciously controlled if desired. Additionally, with the developed method of creating mismatches between oculomotor intentions and reality, it was possible to investigate whether decision-making could be biased by forcing participants to view one alternative while they were naive of such manipulation.

Degree of awareness

Regarding whether people monitor their oculomotor intentions, a few things can be asserted. Whatever neurological activity it is that constitutes oculomotor intentions, be it motor command signals with involvement of activity from the basal ganglia and the frontal eye field, or other types of activities, it does not seem that people in general consciously notice whether these intentions correctly predict the outcome of the eye movements or not, if one accepts the results of this study. Given how the experiment was constructed, it arguably follows that if a person is able to notice that their eye movements affect the movement of the faces, such that it becomes impossible to view one of the faces, it is due to her monitoring the intentions behind her goal-directed eye movements. But only 3 out of 31 participants did notice this explicitly. One could argue that it is likely that it would not be possible to make participants express awareness of 'errors' between their intentions and the outcome of their eye movements after a single or few mismatches. However, the experiment was designed so that if the participants did monitor their oculomotor intentions to some degree, there was a high likelihood of participants positively expressing awareness of the manipulation. This is so because almost half of the decision trials were manipulated, with multiple mismatches occurring each such trial; and the participants did self-control the amount of exposure towards the manipulation, and had to answer many questions pertaining the experience of the decisions in the interview. Therefore, as the majority of participants did not notice mismatches between intentions and outcome according to the analysis of their verbal responses, it is here argued that the participants were not consciously monitoring the mismatches.

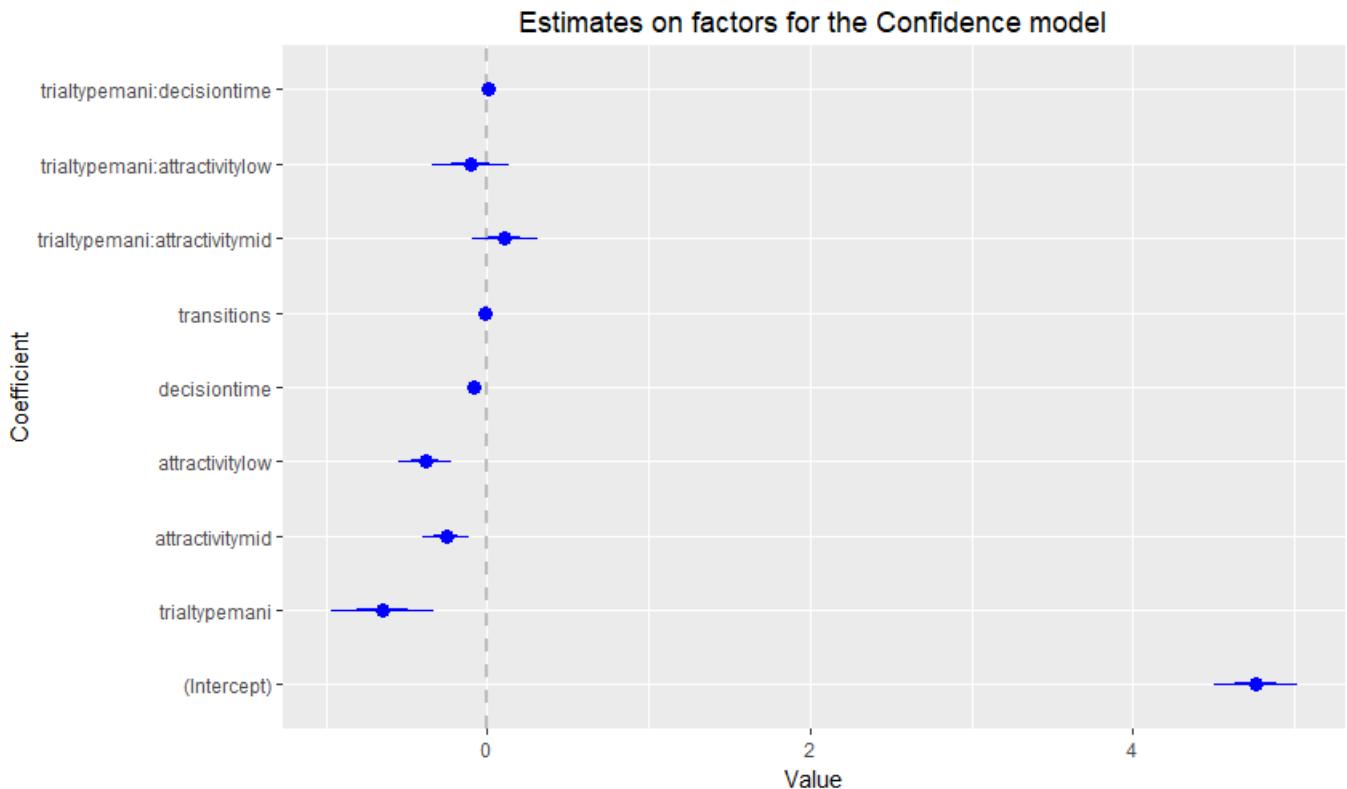


Figure 9. The fixed effects for the coefficient of the model on confidence rating, together with standard errors. The intercept represents the mean when the trial type is non-manipulated, the faces are 'high' in attractivity, and decision time and transitions are set at 0. All significant effects, those from 'transitions' to 'trialtypemani', are negative. 'Transition' refer to the amount of transitions or attempts to transition between one face to the other.

On the other hand, it is highly possible that the manipulation affected participants consciously in other ways, such as their sense of agency, and confidence in their ability to make accurate decisions in a visually demanding task. This is supported by the fact that participants' confidence ratings were significantly lower in manipulated trials. It is also supported from the verbal rapports the participants provided with, as the majority expressed difficulties in perceiving both faces clearly sometimes, which predictably follows from the manipulation.

These results suggest that the symbolic framework, and the goal-directed action models that model cognition similarly, are wrong in proposing that humans have high access to mental states that relate to the active situation, and that specific types of 'conscious-ready' mental events are always necessary in 'higher-order' mental activity, such as deciding what actions to do in a task. The dynamic framework on the other hand supports the idea that 'higher-order' cognitive task, such as deciding which actions to do, can arise naturally as an interaction with the task, with sensory information together with previous learning feeding the decisions without the need of full conscious deliberation, which would require too much attentional resources from an evolutionary point of view.

The predictive processing view, which is an evolvement of the dynamic framework, additionally is able to explain what processes it could be that could lead to a person becoming aware of mismatches between predictions of outcomes and the outcomes themselves, given that baseline human cognition is not symbolic. With this framework, cognition is viewed as primarily predictive in nature, but it does not need consciousness to predict things, which is an explanation for why the majority of the participants in this study did not consciously

become aware of experimental manipulation, only that it was 'difficult to see things'. And as was shown by Ehinger, König and Ossandón (2015), error signals occur after predictive mismatches that arise as a saccade-related change is made, which surprises the visual cortex. The case that 3 participants in this experiment did express awareness of the gaze-contingent manipulation, suggests that they possibly were sensitive to these error signals. Although, it could also be due to increased attention towards eye movements, as they knew their eyes were measured and might have suspected some type of experimental manipulation.

Decision-making

It certainly was the case that the participants' decision-making was affected, even as they were naive to the manipulation, which is supported by the rapports from the participants. Manipulated trials, which forced participants to view the 'back' face, increased the likelihood that participants would choose the 'back' face with about 15% compared to non-manipulated trials, to a 70% chance in choosing that alternative (Fig. 8); but that only applied to trials which included highly attractive faces (highly attractive according to the provided norming tables from the database), which fits with research that show that increased attention towards an alternative during the decision process increases the likelihood of choosing that alternative, but only if it is attractive, with the opposite effect for aversive stimuli (Armel, Beaumel & Rangel, 2008). For manipulated trials with unattractive faces, the likelihood of choosing the 'back' face dropped down by 30%, so that it was only a 40% chance of choosing the only alternative that they had properly

seen, even though the participants had not fixated on the 'front' face at all and hence had a worse estimation of how that face looked like. The participants would likely have become aware of that the faces matched in attractivity during the experiment, so choosing the 'front' face even though they have not seen it properly could not be rationalized with the case that the 'front' face might have been much more attractive by chance.

Effects on recognition and source memory

What is surprising from the data collected is the degree to which participants could distinguish 'blocked' faces. Even though the participants could only view the 'blocked' faces peripherally, their average sensitivity (d') was still above 0. Modeling their recognition rates revealed that the fixed effect from manipulation compared to non-manipulated trials was non-significant, although that includes both 'blocked' and 'forced' faces, which makes it unsurprising that the manipulation factor in itself was non-significant. What can be said is that there was a large and significant difference between front/blocked and back/forced faces in the manipulated condition ('blocked' and 'forced' does not apply to non-manipulated trials). Which means that the faces that were blocked through manipulation were significantly harder to recognize, for all attractivity-levels, although faces with higher attractivity were significantly easier to recognize in all combinations compared to the two other attractivity levels.

The data on participants answering whether they think they chose the presented face in the decision trial it took part in includes an interesting shift in accuracy when changing what position the face had and whether the trial was manipulated or not. On faces participants had claimed they recognized (which were all trials where they were asked whether they think they chose the face in the decision phase), they were on average significantly worse in correctly remembering faces that were positioned 'back' and took part in a non-manipulated trial. This is likely due to the increased bias in looking at the 'back' face in manipulated trials, although that has not been included in the model due to difficulties in modeling that factor in a way that represents it correctly. Then, when shifting over to faces that were positioned in the front, participants were significantly worse at remembering whether they chose the faces in the decision phase if they belonged to a manipulated trial. This is likely due to the fact that the participants general ability to remember the faces that were positioned in the front in manipulated trials was worse. There was no significant effect on attractivity and source memory, although the pattern suggests that there was an effect, but smaller when compared to its effect in the recognition task.

So it does seem to be the case that even if the faces were separated by 5 visual degrees, there was a slight ability to see the 'blocked' faces, and that could possibly have affected the decision-making bias. But if it did affect the bias in decision-making, it would likely have lessened the bias in decision-making for the 'forced' faces. The significant bias that has been observed in this study is then further supported by the fact that there were some visual 'leakages' which made it easier for participants to see blocked faces, and hence made the decisions more equal than the goal of the design which was to make the manipulated decisions almost fully unequal.

Confidence ratings

As previously mentioned, the average participant confidence rating was significantly lower on manipulated trials. Their ratings also decreased significantly when the faces were not in the 'high' attractivity group. But the interaction between manipulation and attractivity was not significant, so it can be said that no matter whether the participants can see both faces clearly, lower attractivity made participants less confident in their decisions. What was also the case was that the longer participants spent on the decision (and how many times they switched or tried to switch between the faces, which are co-linear factors), the less confident the participants became. This was the case even though an interaction effect between decision time and trial type was included. Given that participants controlled the decision time themselves, it seems likely that participants spent more time looking at the faces if they were uncertain of which face to choose, and that more observation-time might not have helped much as the task was visually demanding.

The problem with having confidence ratings saying anything about the participants' sense of agency is that confidence ratings are likely affected by many factors—one's sense of agency is only one of those. But no developed measures for sense of agency could be applied in this task, as the few implicit measures that exist, sensory attenuation and temporal binding (as opposed to explicit measures where participants are asked questions pertaining to their sense of agency, in for example questionnaires; see Dewey & Knoblich, 2014), require controlling the experimental events strictly, and involve participants doing specific individual actions. In this study, the specific actions we were interested in were extremely frequent, and in the minds' of the participants the eye movements they made were uninteresting (as long as they did not notice how the manipulation occurred), and that had to be the case as we needed naive participants. But given that the confidence ratings did significantly decrease, it makes sense that the participants' sense of agency would be lowered in an experiment like this, although it cannot be proven.

Limitations

The primary limitation with the experiment regarded the post-test interview. There are some obvious limitations in interviewing participants and then analyzing their responses as indications of their degree of awareness. It could fully be the case that participants were not motivated, or had not the concepts and language to express the feeling of failure with regards to oculomotor intentions (as the concept of oculomotor intentions is not common in public discourse). There is also the balance between activating participants to fully exhaust in language their subjective experience, so that we can determine whether they did or did not become aware of failures of oculomotor intentions, and on the other hand not to provide leading questions that participants would use to create a confabulation to fit the question. Even explicitly stating the manipulation at the end and analyzing the participants' responses can not be a secure measure of participant awareness, as the participants' responses can be confabulated even here to fit the demand. However, most participants did not express that they noticed mismatches between intentions and the result of their eye movements, even when told about the manipulation in the debriefing, and no participant did express that they monitored their oculomotor intentions directly. The participants who noticed the manipulation also expressed only noticing how the

system functioned and the role of their eye movements in the system, so it still is not exactly clear what these participants really experienced. It is for instance not clear if they actually thought about *intentions*, or just the eye movement effect in the movement of the faces in the experiment.

Summary and suggestions for further research

The results from the study suggest that oculomotor intentions are not explicitly monitored by most people. If some people do, it might have to do with sensitivity towards error signals that likely occur after gaze-contingent mismatches. This supports the idea that we humans generally do not know ourselves as good as we like to think, and that we do not have full introspective access to all our 'higher-level' mental activity, such as deciding which action to express physically. But the study does not exclude that oculomotor intentions are completely unavailable to conscious deliberation, both before and after an intentional action has been made.

Given that most participants did not notice the manipulation even when almost half of all trials were manipulated, and that they controlled the duration of exposure to the manipulation, a further step would be to push an experimental design of similar nature to the limit, to determine how much would be required for participants to spontaneously notice the mismatch between their oculomotor intentions and the result of their eye movements. A sophisticated way of doing this is with the use of virtual reality (VR). It could be possible to situate participants in a virtual environment, where the whole environment itself, or parts of it, are gaze-contingent, such that participants could be tested in more a 'natural' environment (more natural than sitting still in front of a monitor and making eye movements towards faces that rapidly jump on a screen), where intentions to move one's gaze towards objects would be easier to define and measure, than the small and frequent intentional eye movements that occurred in this experiment. In more 'natural' situations, it might turn out that people generally do monitor such mismatches between intentions and outcome, or monitor other informative perceptual disturbances that could occur when the brain by habit predicts the visual information before the movement, but receives something completely unexpected after the movement.

5 Conclusion

This study provides evidence that humans do not generally monitor their oculomotor intentions. It was also found that decision-making could be biased by blocking participants in viewing one of the alternatives in a forced-choice task, without participants knowing that they were manipulated in that way. This suggests that models that assume that intentions are a necessary step prior to a goal-directed action need to be revised.

6 Acknowledgements

I would like to thank Petter Johansson and Lars Hall for their interesting and creative ideas, and their support and feedback in developing the experimental design, and I would like to thank Philip Pärnamets for providing incredible technical and statistical support, whom without I could not have achieved the results I got.

7 References

- Ajzen, I. (1991). The theory of planned behavior. *Organizational behavior and human decision processes*, 50, 179-211.
- Armel, K. C., Beaumel, A., and Rangel, A. (2008). Biasing simple choices by manipulating relative visual attention. *Judgment and Decision Making*, 3, 396.
- Bates, B., Maechler, M., Bolker, B. and Walker S. (2015). Fitting Linear Mixed-Effects Models Using lme4. *Journal of Statistical Software*, 67, 1-48.
- Beeler, G. W. (1967). Visual threshold changes resulting from spontaneous saccadic eye movements. *Vision research*, 7, 769-775.
- Blakemore, S. J., Wolpert, D., and Frith, C. (2000). Why can't you tickle yourself?. *Neuroreport*, 11, R11-R16.
- Bridgeman, B., Hendry, D., and Stark, L. (1975). Failure to detect displacement of the visual world during saccadic eye movements. *Vision research*, 15, 719-722.
- Buschman, T. J., and Miller, E. K. (2014). Goal-direction and top-down control. *Phil. Trans. R. Soc. B*, 369, 20130471.
- Carrasco, M. (2011). Visual attention: The past 25 years. *Vision research*, 51, 1484-1525.
- Carruthers, P. (2011). *The opacity of mind: an integrative theory of self-knowledge*, Oxford University Press, Oxford.
- Clark, A. (2015). *Surfing uncertainty: Prediction, action, and the embodied mind*, Oxford University Press, Oxford.
- Clarke, A. D., Mahon, A., Irvine, A., and Hunt, A. R. (2017). People are unable to recognize or report on their own eye movements. *The Quarterly Journal of Experimental Psychology*, 70, 2251-2270.
- Della Sala, S., and Marchetti, C. (2005). Anarchic hand. In F. Hans-Joachim, M. Jeannerod, M. Hallett, and R. Leiguarda (Eds.), *Higher-order motor disorders: from neuroanatomy and neurobiology to clinical neurology* (pp. 291–301). Oxford University Press, Oxford.
- Dennett, D. C. (1991). Real patterns. *The journal of Philosophy*, 88, 27-51.
- Deubel, H. (2004). Localization of targets across saccades: Role of landmark objects. *Visual Cognition*, 11, 173-202.
- Deubel, H., and Schneider, W. X. (1996). Saccade target selection and object recognition: Evidence for a common attentional mechanism. *Vision research*, 36, 1827-1837.
- Dewey, J. A., and Knoblich, G. (2014). Do implicit and explicit measures of the sense of agency measure the same thing?. *PloS one*, 9, e110118.
- Dickinson, A., and Balleine, B. (1994). Motivational control of goal-directed action. *Animal Learning & Behavior*, 22, 1-18.
- Ehinger, B. V., König, P., and Ossandón, J. P. (2015). Predictions of visual content across eye movements and their modulation by inferred information. *Journal of Neuroscience*, 35, 7403-7413.
- Haggard, P. (2017). Sense of agency in the human brain. *Nature Reviews Neuroscience*, 18, 196-207.
- Hommel, B. (2015). Action control and the sense of agency. In Haggard, P. and Eitam, B. (Eds.), *The sense of agency* (pp. 307-326). Oxford Scholarship Online.
- Johansson, P., Hall, L., Sikström, S., and Olsson, A. (2005). Failure to detect mismatches between intention and outcome in a simple decision task. *Science*, 310, 116-119.
- Fodor, J. A. (1985). Fodor's guide to mental representation: The intelligent auntie's vade-mecum. *Mind*, 94, 76-100.

- Friston, K. (2012). Prediction, perception and agency. *International Journal of Psychophysiology*, 83, 248-252.
- Gilron, R., Simon, S. and Mukamel, R. (2015). Neural correlates of intention. In Haggard, P. and Eitam, B. (Eds.), *The sense of agency* (pp. 95-114). Oxford Scholarship Online.
- Girard, B., and Berthoz, A. (2005). From brainstem to cortex: computational models of saccade generation circuitry. *Progress in neurobiology*, 77, 215-251.
- Guo, K., Liu, C. H., and Roebuck, H. (2011). I know you are beautiful even without looking at you: discrimination of facial beauty in peripheral vision. *Perception*, 40, 191-195.
- Hosoya, T., Baccus, S. A., and Meister, M. (2005). Dynamic predictive coding by the retina. *Nature*, 436, 71-77.
- Hunziker, H.W. (2006). *Im Auge des Lesers: foveale und periphere Wahrnehmung—vom Buchstabieren zur Lese-freude* [In the eye of the reader: foveal and peripheral perception—from letter recognition to the joy of reading]. Transmedia Staubli, Zurich.
- Kowler, E. (2011). Eye movements: The past 25 years. *Vision research*, 51, 1457-1483.
- Krajbich, I., Armel, C., and Rangel, A. (2010). Visual fixations and the computation and comparison of value in simple choice. *Nature neuroscience*, 13, 1292-1298.
- Laird, J. E., Newell, A., and Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial intelligence*, 33, 1-64.
- Lappi, O. (2016). Eye movements in the wild: oculomotor control, gaze behavior & frames of reference. *Neuroscience & Biobehavioral Reviews*, 69, 49-68.
- Lisman, J. (2014). Two-phase model of the basal ganglia: implications for discontinuous control of the motor system. *Phil. Trans. R. Soc. B*, 369, 20130489.
- Lupyan, G., and Clark, A. (2015). Words and the world: Predictive coding and the language-perception-cognition interface. *Current Directions in Psychological Science*, 24, 279-284.
- Ma, D. S., Correll, J., and Wittenbrink, B. (2015). The Chicago face database: A free stimulus set of faces and norming data. *Behavior research methods*, 47, 1122-1135.
- Melcher, D., and Colby, C. L. (2008). Trans-saccadic perception. *Trends in cognitive sciences*, 12, 466-473.
- Montoya, R. M., Horton, R. S., Vevea, J. L., Citkowicz, M., and Lauber, E. A. (2017). A re-examination of the mere exposure effect: The influence of repeated exposure on recognition, familiarity, and liking. *Psychological Bulletin*, 143, 459-498.
- Newell, A., and Simon, H. A. (1976). Computer science as empirical inquiry: Symbols and search. *Communications of the ACM*, 19, 113-126.
- Nisbett, R. E., and Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological review*, 84, 231.
- Nittono, H., and Wada, Y. (2009). Gaze shifts do not affect preference judgments of graphic patterns. *Perceptual and motor skills*, 109, 79-94.
- O'regan, J. K. (1992). Solving the "real" mysteries of visual perception: the world as an outside memory. *Canadian Journal of Psychology/Revue canadienne de psychologie*, 46, 461.
- Peirce, J.W. (2007). PsychoPy - Psychophysics software in Python. *Journal of Neuroscience Methods*, 162, 8-13.
- Prescott, T. J. (2008). Action selection. *Scholarpedia*, 3, 2705.
- Pärnamets, P., Johansson, P., Hall, L., Balkenius, C., Spivey, M. J., and Richardson, D. C. (2015). Biasing moral decisions by exploiting the dynamics of eye gaze. *Proceedings of the National Academy of Sciences*, 112, 4170-4175.
- Redgrave, P., Rodriguez, M., Smith, Y., Rodriguez-Oroz, M. C., Lehericy, S., Bergman, H., Agid, Y., DeLong, M. R. and Obeso, J. A. (2010). Goal-directed and habitual control in the basal ganglia: implications for Parkinson's disease. *Nature Reviews Neuroscience*, 11, 760-772.
- Scherberger, H., Jarvis, M. R., and Andersen, R. A. (2005). Cortical local field potential encodes movement intentions in the posterior parietal cortex. *Neuron*, 46, 347-354.
- Shimojo, S., Simion, C., Shimojo, E., and Scheier, C. (2003). Gaze bias both reflects and influences preference. *Nature neuroscience*, 6, 1317-1322.
- Sommer, M. A., and Wurtz, R. H. (2008). Brain circuits for the internal monitoring of movements. *Annual Review of Neuroscience*, 31, 317-338.
- Sparks, D. L. (2002). The brainstem control of saccadic eye movements. *Nature Reviews Neuroscience*, 3, 952-964.
- Thakkar, K. N., Diwadkar, V. A., and Rolfs, M. (2017). Oculomotor Prediction: A Window into the Psychotic Mind. *Trends in Cognitive Sciences*, 5, 344-356.
- Todd, J., Michie, P. T., Schall, U., Ward, P. B., and Catts, S. V. (2012). Mismatch negativity (MMN) reduction in schizophrenia—Impaired prediction-error generation, estimation or salience?. *International Journal of Psychophysiology*, 83, 222-231.
- Triesch, J., Ballard, D. H., Hayhoe, M. M., and Sullivan, B. T. (2003). What you see is what you need. *Journal of vision*, 3, 9-9.
- Van Gelder, T. (1995). What might cognition be, if not computation? *The Journal of Philosophy*, 92, 345-381.
- Verschure, P. F., Pennartz, C. M., and Pezzulo, G. (2014). The why, what, where, when and how of goal-directed choice: neuronal and computational principles. *Phil. Trans. R. Soc. B*, 369, 20130483.
- Wojnowicz, M. T., Ferguson, M. J., Dale, R., and Spivey, M. J. (2009). The self-organization of explicit attitudes. *Psychological Science*, 20, 1428-1435.
- Yarbus, A. L. (1967). *Eye Movements and Vision*, Plenum Press, New York.