Running head: Consciousness During Learning

Learning Rapidly about the Relevance of Visual Cues Requires Conscious Awareness

Eoin Travers, Chris D. Frith, Nicholas Shea

Institute of Philosophy, School of Advanced Study, University of London

Author Note

1

Abstract

Humans have been shown capable of performing many cognitive tasks using information of which they are not consciously aware. This raises questions about what role consciousness actually plays in cognition. Here, we explored whether participants can learn cue-target contingencies in an attentional learning task when the cues were presented below the level of conscious awareness, and how this differs from learning about conscious cues. Participants' manual (Experiment 1) and saccadic (Experiment 2) response speeds were influenced by both conscious and unconscious cues. However, participants were only able to adapt to reversals of the cue-target contingencies (Experiment 1) or changes in the reliability of the cues (Experiment 2) when consciously aware of the cues. Therefore, although visual cues can be processed unconsciously, learning about cues over a few trials requires conscious awareness of them. Finally, we discuss implications for cognitive theories of consciousness.

Introduction

Much of our mental life is conscious. We not only perceive our environment, engage cognitive processes, and execute movements, but we are often aware that we are doing so. These states and processing steps figure in our subjective experience. However, a large body of evidence (e.g. Boy & Sumner, 2010; Pessiglione et al., 2008; van Gaal, Ridderinkhof, van den Wildenberg, & Lamme, 2009; Wenke, Fleming, & Haggard, 2010) shows that many of these activities can be performed without subjective awareness: we can perceive stimuli, make decisions, and act, without being conscious of doing so.

This raises the question: what role does consciousness play in cognitive processes? According to some accounts (e.g. Baumeister, Masicampo, & Vohs, 2011; Huxley, 1874; Pockett, Banks, & Gallagher, 2009) consciousness is an epiphenomenon — a causally inert by-product which makes no difference to the way cognition unfolds and actions are performed. According to others, consciousness plays a direct causal role in high-level cognition, for instance by facilitating the broadcast of information between different neural systems (Dehaene & Changeux, 2011; Dehaene, Kerszberg, & Changeux, 1998) or the integration of information from different sources (Tononi, 2004), monitoring of lower-level cognitive processes (Norman & Shallice, 1986) by allowing mental states to be communicated to others (Frith, 2010) or allowing us to think counterfactually about our own decisions (Byrne, 2016; Frith & Metzinger, 2016).

One productive line of inquiry has been to identify ways in which information processing takes place differently when representations are conscious, compared to the operations that can be performed on information that is represented unconsciously (see Dehaene, Charles, King, & Marti, 2014 for a review). The strongest such evidence comes from paradigms where stimuli are rendered unconscious by a masking procedure (see Kouider & Dehaene, 2007). In this way, a surprising number of processes have been found to operate on non-conscious representations. For instance, humans can process simple directional cues (Eimer & Schlaghecken, 1998), recruit executive processes (van Gaal et al., 2009; Wokke, van Gaal, Scholte, Ridderinkhof, & Lamme, 2011), respond viscerally to emotional stimuli (Kiss & Eimer, 2008; Winkielman & Berridge, 2004), and partially integrate the meanings of multiple words (van Gaal et al., 2014), without conscious awareness of the stimuli.

There may be a qualitative difference between the kinds of tasks that can be performed without consciousness and those that require consciousness. Bekinschtein et al. (2009), for instance, showed that participants' ERPs reflected detection of local violations in temporal structure that participants were consciously unaware of, but that this detection only occurred for global violations when participants were conscious of them. However, given the increasing number of processes that

3

have been shown to occur in the absence of conscious awareness, it is no longer certain that any phenomena exist that strictly require consciousness (Dehaene et al., 2014; Shea & Frith, 2016). An alternative hypothesis is that a representation's being conscious serves to facilitate a range of cognitive processes that can be performed on it, each of which could individually be performed in some fashion in the absence of consciousness (Shea & Frith, 2016).

In investigating which tasks consciousness is necessary for or facilitates, a capacity that has received relatively little attention is learning: does the way we update our beliefs about the world in response to new information depend on whether we are conscious of this information? We note that this article focuses on unconscious learning — learning from information of which we are not consciously aware — and not implicit learning — learning without being explicitly aware of what has been learned. There has been considerable work on implicit learning, using both artificial grammar learning and classical conditioning, and the topic remains somewhat contentious (Clark, 1998; Lovibond & Shanks, 2002; Vadillo, Konstantinidis, & Shanks, 2015). It is also worth noting that unconscious learning is not the same as learning without paying attention to the information being learned (e.g. Seitz & Watanabe, 2003)

Of course, there are many kinds of learning, and in the current article we focus on contingency learning, or classical conditioning. At its most basic, this is simply learning that some stimuli tend to occur together or in close succession, and how strong these contingencies are (e.g. Fanselow & Poulos, 2005). There is good evidence that participants can do this in quite sophisticated ways. Vossel et al. (2014), for instance, presented an experiment using a version of the attentional cueing task (Posner, 1980) where arrow cues (left or right) predicted the location of target stimuli in a detection task, and showed that participants flexibly adjusted their reliance on the cues – the extent to which they were faster to respond when the cues were valid, and slower when they were invalid – depending on both the proportion of cues that were valid, and the extent to which this proportion was subject to change (volatility). There is evidence that participants can achieve other forms of learning about stimuli they do not consciously see. For example, (Pessiglione et al., 2008) rewarded participants for pressing a button in response to one stimulus, and punished them for doing so in response to another, and found that participants could learn to respond more often to the rewarded stimulus even when the stimuli were rendered unconscious by masking, indicating that operant conditioning can occur for unconscious stimuli (see also Atas, Faivre, Timmermans, Cleeremans, & Kouider, 2014). There is mixed evidence, however, that contingency learning can also occur without conscious awareness. Interestingly, this evidence comes from work exploring conflict adaptation effects (see Ansorge, Kunde, & Kiefer, 2014; Desender & Van Den Bussche, 2012, for recent reviews). Conflict adaptation occurs on classical conflict tasks (e.g. Eriksen & Eriksen, 1974; Simon & Wolf, 1963; Stroop, 1935) when the influence of conflicting stimuli changes over time. The conflict effect – the difference between responses to congruent and incongruent trials – is modulated

by participants' previous experiences of conflict. Conflict effects are both reduced on trials immediately following conflict trials (Gratton, Coles, & Donchin, 1992), and reduced on blocks where there are many incongruent trials (Jacoby, Lindsay, & Hessels, 2003): the more often a cue conflicts with the target, the less it will be processed. The question of whether or not conflict adaptation occurs for unconscious stimuli is important, as cognitive control processes – those needed to overcome conflict – are traditionally seen as exclusively conscious in nature (Dehaene & Naccache, 2001), and evidence for the unconscious activation of these processes has implications for many theories of consciousness. Our interest in these phenomena, however, is more general, in that we are interested in unconscious learning more broadly, rather than the deployment of control processes.

It is also worth noting that learning or adaptation here can take a number of forms. For instance, participants could learn to what extent a cue, including a previously neutral one, is predictive of a particular target (simple contingency learning), learn that a cue elicits an unhelpful automatic response, and so its processing should be inhibited (conflict adaptation), or learn that a cue that was predictive of one target is now predictive of the other (reversal learning). Moreover, our current interest is in *rapid* learning. It is commonly accepted (McClelland, Rumelhart, et al, 1986) that even simple neural architectures can adapt over time by Hebbian learning, and recent work (Gagliano, Vyazovskiy, Borbély, Grimonprez, & Depczynski, 2016) has even claimed to demonstrate slow associative learning in plants. Our research question, therefore, is how conscious and unconscious learning differ, and if rapid learning can be achieved for unconscious stimuli. We would also point out that in conflict adaptation the cue automatically affects behaviour, and so is referred to as a prime. For simplicity, we will use the term "cue" throughout this paper.

It is well established that participants adapt to the prevalence of incongruent trials on standard conflict tasks, even when they are not explicitly aware of the actual proportion of conflict trials (Blais, Harris, Guerrero, & Bunge, 2012; Crump, Gong, & Milliken, 2006; Ridderinkhof, 2002). It has also been found that when a cue is presented subliminally, participants can use conscious knowledge about the likelihood of it being incongruent – obtained either from learning about consciously-presented cues, or from explicit instruction – to modulate how much they allow the cue to influence their behaviour (Jiang, van Gaal, Bailey, Chen, & Zhang, 2013; Reuss, Pohl, Kiesel, & Kunde, 2011). More importantly for the topic at hand, similar effects have been reported in studies using only subliminally masked stimuli, where participants were not informed of the validity of the cues. In a number of papers, Bodner and colleagues (Bodner & Dypvik, 2005; Bodner & Masson, 2003, 2004, 2001; Bodner, Masson, & Richard, 2006) report that lexical or numerical priming effects are reduced for participants who experienced a large number incongruent trials, compared to those who only experienced a few. Similarly, Klapp (2007) presented participants with a task (Eimer & Schlaghecken, 1998) where target arrows were preceded by primes that were either identical or pointed in the opposite direction, and found that the influence of unconsciously presented primes

varied according to what proportion of primes were identical to the target. A serious limitation of these studies from our point of view, however, is that the proportion of congruent trials was manipulated only between participants, making it difficult to assess each participant's actual learning. A within-participants manipulation was used by Bodner and Stalinski (2008), who also manipulated cognitive load between participants, and, surprisingly, found adaptation effects for participants under cognitive load, but not those under normal conditions. This suggests that the effects found in between-participant designs may not be due to actual learning over the course of the experiment. Jaśkowski, Skalska, and Verleger (2003, Experiment 3) do report apparent within-participant unconscious conflict adaptation, but tested only 9 participants, and do not report a number of key details.

Another substantial issue for many of these previous studies is that control of the subjective visibility of the primes was poor: in almost every case, a considerable portion of participants reported being able to at least somewhat see the masked cues, and although analyses are reported with these participants excluded, Vadillo et al. (2015) have demonstrated that this approach is not adequate to demonstrate unconscious processing. A small number of studies (i.e. Bodner & Mulji, 2010; Jaśkowski et al., 2003; Klapp, 2007) have presented participants with a forced-choice prime identification test at the end of their experiment, and tested the correlation between participants' performance on the identification test and their priming effect, although issues remain with the statistical power of some these tests. Perhaps most importantly, however, to our knowledge no previous studies have investigated the relationship between participants' ability to identify the masked stimuli and actual conflict adaptation effects, the phenomenon of interest, leaving it unclear if the adaptation effects were due to a subset of participants who were somewhat aware of the primes. This question is crucial for demonstrating unconscious learning, and we will return to it later in this paper.

Recently, studies of unconscious conflict adaptation have begun to focus on adaptation to the proportion of congruent trials within a specific context, where the context changed trial-by-trial, rather than across a whole block of trials. This shift has occurred largely in response to concerns (e.g. Desender & Van Den Bussche, 2012; Jaśkowski et al., 2003; Reuss, Desender, Kiesel, & Kunde, 2014) that even when primes themselves are processed unconsciously, their blockwise validity may be inferred from their subsequent metacognitive effects: participants may become aware that the mostly-incongruent condition is unusually difficult, either because they notice they are generally slower to respond or make more errors, or because they otherwise infer it from their own internal confidence or sense of difficulty. In the context-specific paradigm, in one context most cues are congruent, and in the other most are incongruent. Context can be signalled, for instance, by the type of mask used (Panadero, Castellanos, & Tudela, 2015), features of the cue, or features of the target itself (Reuss et al., 2014). Whether or not context-specific conflict adaptation occurs for unconscious stimuli remains a somewhat open question. Some studies have reported positive results (e.g. Reuss et al., 2014), some negative (Heinemann, Kunde, & Kiesel, 2009), and some have shown inconsistencies

even within a single paper (Schouppe, de Ferrerre, Van Opstal, Braem, & Notebaert, 2014). Additionally, it appears that effects on response times and error rates may be somewhat dissociated, as a number of papers have reported unconscious effects for one, but not the other. Despite this, context-specific adaptation remains a key topic in cognitive control, largely because the paradigm eliminates a number of alternative explanations that could be offered for apparent blockwise learning. In the current work, however, we focus on learning and adaptation between experimental blocks. We do so for a number of reasons. First, a number of issues with the studies discussed above remain unresolved, and the existence of unconscious learning more broadly remains contentious. It is worthwhile, therefore, to further investigate the existence of blockwise adaptation effects while addressing the limitations of previous work, in particular the poor control of cue visibility, and the failure to test the relationship between cue visibility and adaptation effects. Second, although the metacognitive factors highlighted by Desender & Van Den Bussche (2012) and others are indeed problematic for accounts of unconscious cognitive control, our interest here is in learning more generally. Our question, therefore, is whether or not learning occurs for unconscious stimuli when the issues raised above are addressed, and if so how it differs from conscious learning, not the route by which this might come about. Additionally, blockwise learning is the more commonly-studied phenomenon outside of the question of conscious awareness. A number of studies have shown that human learning can be remarkably adaptive: people construct rich mental models of their environments, and update these models in an approximately Bayesian manner in response to new evidence (Mathys, Daunizeau, Friston, & Stephan, 2011; Tenenbaum, Kemp, Griffiths, & Goodman, 2011). If learning does occur for unconscious cues, it is natural to ask how this differs from normal conscious learning. Third, given our aim here is to be more stringent than previous work with relation to our control of stimulus visibility, it is worth being more liberal in other regards. In other words, if we do find evidence against unconscious learning, it is preferable to have done so using a paradigm where rapid learning unambiguously occurs for conscious stimuli.

To compare rapid learning about conscious information with our capacity for learning about unconsciously represented information, we adapted a classic attentional cuing task (Posner, 1980; Vossel et al., 2014). In this, participants are required to respond quickly to targets on either side of the screen. A cue is shown before these targets, and these cues can be used as probabilistic predictors of the location of the target, left or right, facilitating rapid responses. If participants successfully learn the contingencies between the cue and the target location, they should allocate their attention accordingly, and so respond faster, with fewer errors, on trials that follow these contingencies, and respond more slowly, and commit more errors, when these expectations are violated (Vossel et al., 2014). In the standard version of this paradigm, arrows are used as cues, and participants almost certainly begin with the assumption that arrows are positive predictors of the location of the target, and so their automatic response will be to expect the target on the indicated side. As noted above, a

number of different factors may be in play here. Participants should be initially driven by their automatic response to the cues, and, if learning occurs, also modulate these responses by changing their weighting of the cues, or even learning to expect a different target in response to them.

Previous work (McCormick, 1997; Webb, Kean, & Graziano, 2016) has shown that attention on this task can be unconsciously captured by the abrupt onset of peripheral cues: attention can be drawn towards a pre-target stimulus on either side of the screen, even when participants are not aware of this stimulus. However, orienting towards a novel stimulus is a relatively low-level capacity, known to be present from birth, as well as in non-human animals (see Sokolov, 1963). Processing centrally located cues, on the other hand — arrows pointing towards one or other side of the screen — is in principle a considerably more complex operation, requiring participants to process the meaning of the cue and orient their attention accordingly. There is extensive evidence (Boy & Sumner, 2010; Eimer & Schlaghecken, 1998, 2003), however, that subliminal central cues can prime both attention and responses. It is worthwhile to note out that in these experiments the cue is typically equally likely to be congruent as incongruent, and so these effects appears to be the result of an already-acquired automatic response to the cue itself, rather than learning over the course of the experiment. Furthermore, in using this paradigm, we make no attempt to differentiate between the effect of the cue on attention (facilitating faster detection of the target at the cued location) and the direct priming of left- or right-hand responses. Previous work (Al-Janabi & Finkbeiner, 2014) suggests the effect of masked cues may primarily due to response priming, rather than attentional priming. As noted above, our desire is to allow every opportunity for unconscious rapid learning to occur, and so we have opted for a paradigm with multiple possible routes to priming.

In the current work, we ask if participants can learn about such cue-target contingencies, and so adjust how they respond to cues, without being consciously aware of the cues themselves. In Experiment 1, we follow a procedure similar to that used by Jiang et al. (2013), and present participants with both cues that point to where the target will be 75% of the time and, in a separate block, cues that point in the opposite direction 75% of the time. Unlike Jiang et al. (2013), however, we do not mix conscious and unconscious cues within a block, and so any adaptation that occurs must be due to learning from the unconscious cues.

In Experiment 2, we investigate a more subtle form of learning, and, following Vossel et al. (2014), present participants with blocks where cues point towards where the target will be 50%, 69%, or 88% of the time. By presenting the cues either consciously or unconsciously in different blocks (Experiment 1) or between participants (Experiment 2), we can isolate what aspects of cue utilisation and learning only occur for cues that are represented consciously, and what aspects can occur for unconscious cues.

Experiment **1**

In Experiment 1, we presented participants with double arrow cues (« or ») that pointed towards where the target would be (i.e. cues that were congruent with the target) on 75% of trials in one block, and 25% of trials in another block. In the latter case the cue was a good (75%) predictor that the target would appear on the opposite side to where the arrow is pointing (i.e. in the incongruent direction). If participants utilise these cues, their response times should differ between trials where the target appears on the side pointed to by the cue and trials where it appears the opposite side. If they learn about the cues, this difference should vary between the blocks where most cues pointed to the same side and the blocks where most pointed to the opposite side.

Method

Participants

Twenty-eight participants were recruited and took part in the experiment in a quiet area of the University of London Students' Union (mean age = 26.3 years, SD = 10.1). All reported having normal or corrected-to-normal vision. Each received £5 for taking part in the experiment.

Design

The main experiment consisted of 400 trials. Cues were presented consciously for one half of the experiment, and unconsciously for the other half, with the ordering counterbalanced across participants. Each half of the experiment was split into 2 blocks of 100 trials each: a congruent block, where the cues pointed to where the target was going to appear 75% of the time and the opposite direction 25%, and an incongruent block, where they pointed to where the target was going to appear 25% of the time and the opposite direction 75%. This was also counterbalanced across participants, so there were four counterbalanced conditions in total: conscious cues first, congruent block first; conscious cues first, incongruent block first; unconscious cues first, congruent block first; and unconscious cues first, incongruent block first. Participants were not made aware of the transitions between blocks.

After the main experiment, participants completed a visibility check, where the cues were presented as they were in the experiment, and participants were asked to indicate which way the cues pointed. Each participant did so 24 times for cues intended to be conscious, and then 24 times for cues intended to be unconscious. After each block of 24 trials, participants indicated to what extent they

saw the cues on the previous trials, on a 9-point slider ranging from "Completely guessing" (left) to "Completely saw them" (right).

Stimuli and Procedure

Participants were seated approximately 60 cm from a laptop computer which controlled the experiment. All stimuli were presented in black on a white background, using OpenSesame experiment builder and the PsychoPy back end. Responses were collected using a Cedrus response box, and participants rested the fingers of their left and right hands on the respective left and right buttons.

Unless otherwise stated, all stimuli subtended 1.5° visual angle. The fixation symbol (+) was located in the centre of the screen. This was presented for 750 msec before each trial, with small dots located 3° above and below it at the points where the cues could appear. The fixation thickened slightly for the final 100 msec to warn participants that the trial was about to begin. It then disappeared, to be replaced by the cue. As cues, we used left and right pointing double arrows (« and ») subtending 1.5° vertically and 3° horizontally. These cues were randomly presented at 3° above or below fixation in order to avoid the unconscious inhibition effects reported by Schlaghecken and Eimer (2000) for centrally positioned cues. These were shown for 400 (conscious condition) or 33 (unconscious condition) msec before being hidden by the mask, consisting of both right and leftward cues superimposed, shown for 33 msec. Finally, the target was presented: an X symbol offset to 7° left or right of the fixation. Participants were then required to respond as quickly as possible by pressing the corresponding button with their left or right hands. The stimuli for the post-test check were identical, except that instead of the target participants saw text reading "Left or Right", presented centrally, and were asked to indicate which way the cue had pointed.
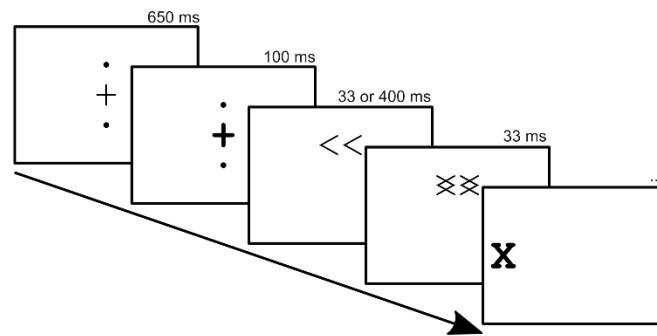
Figure 1. Task from Experiment 1 (not to scale). Double arrow cues pointed either left or right, and where shown slightly above or below fixation, before being covered by a pattern mask consisting of both cues superimposed for 33 msec, followed by the onset of the target on the left or right of the screen. Unconscious cues were shown for 33 msec, and conscious cues for 400. In Experiment 2, unconscious cues were shown for 50 msec, and the target X was replaced by a Gabor patch.

Results

Data Exclusions

We excluded data from one participant who performed significantly below chance, and one did not differ from chance, for the 400 msec cues on the visibility check. Over the remaining with 10,400 trials, we excluded 20 trials with RTs greater than 1.5 seconds, and 120 additional trials RTs more than 3 SDs above that participant's mean.

Errors

Error rates, as a function of cue direction, block, block order, and cue visibility, are shown in Figure 2. Due to the extremely low error rate (3.2% overall), we do not report a statistical analysis of these errors. However, errors were generally more frequent when the cue pointed to the opposite side to the target. This was the case for both conscious and unconscious cues, but appears to only hold for participants who completed the congruent block first. On conscious trials, these participants were more affected by the cue in the congruent block, as expected. On unconscious trials, unexpectedly, the influence of the cues was greater in the incongruent block.

Response times

For analysis, we calculated the *cueing effect*: the difference between mean response times on trials where the cue pointed to where the target would appear and response times when the cue pointed in the opposite direction – or how much slower participants were to respond when the cue pointed in the opposite direction – for each participant, in each block (Figure 3). We subjected this measure to an initial visibility (conscious or unconscious) x block (mostly congruent or mostly incongruent) x block order (mostly congruent first or mostly incongruent first) mixed ANOVA, with block order as a between-participants factor. Consistent results were obtained when using the percentage change in RT (that is, the difference in log RT), and when conducting a full cue (same or opposite side) x visibility x block x block order ANOVA.

We found a significant main effect of block, $F(1,24) = 8.577$, $p = .007$, $\eta^2 = 0.078$, such that the cueing effect was greater on blocks where most cues were congruent. However, this was modulated by significant interactions between visibility and block, $F(1,24) = 11.839$, $p = .002$, $\eta^2 = 0.100$, and between visibility and block order, $F(1,24) = 4.585$, $p = .043$, $\eta^2 = 0.059$. Therefore, we conducted separate analyses for the conscious and unconscious trials, fitting block (mostly congruent or mostly incongruent) x block order (mostly congruent first or mostly incongruent first) mixed ANOVAs to each.
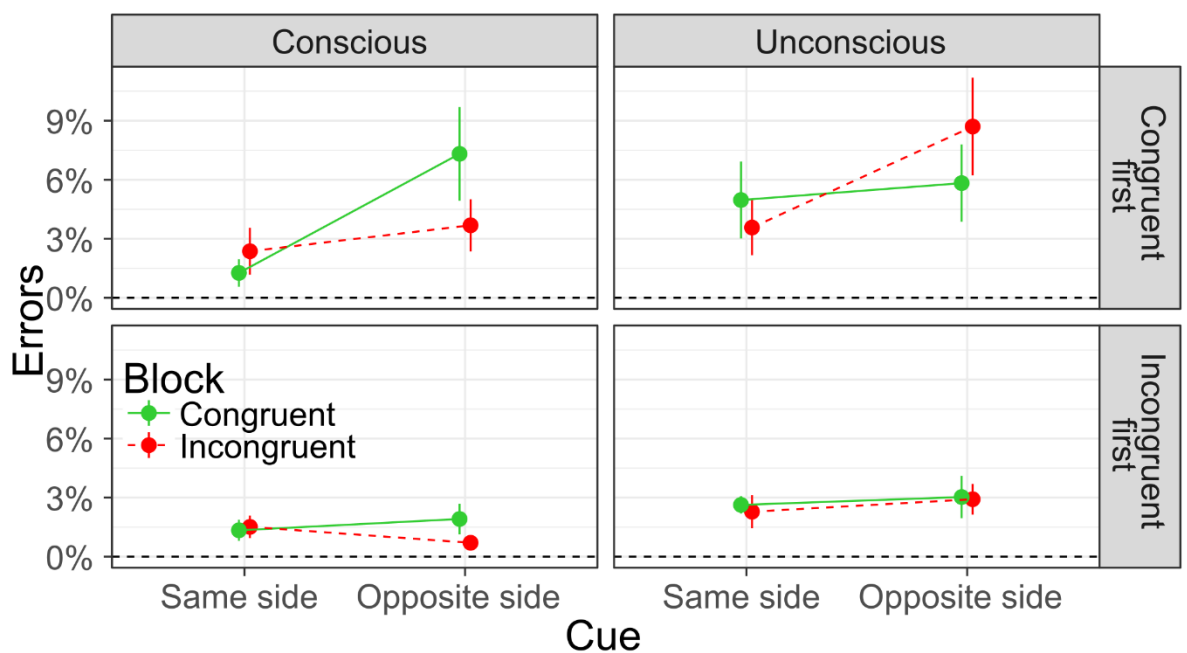

Figure 2. Errors, as a function of cue direction, block, block order, and cue visibility, Experiment 1.
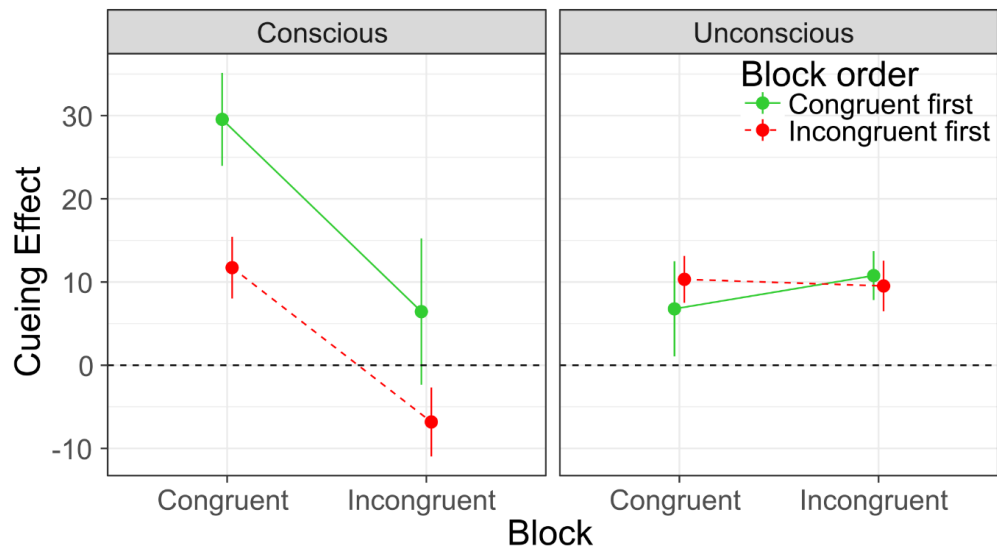
Figure 3. Cueing effect on response times, as a function of cue visibility, block, and block order, Experiment 1

For the conscious cues, this revealed significant main effects of block, $F(1,24) = 17.764$, $p < .001$, $\eta^2 = 0.217$, BF_{alt} (Bayes Factor in favour of the alternative hypothesis) = 52.4, indicating that the cueing effect was greater on blocks where most cues were congruent, and of block order, $F(1,24) = 6.013$, $p = .022$, $\eta^2 = 0.135$, $BF_{alt} = 5.8$, indicating that cueing effects were greater overall for participants who completed the mostly congruent block first. There was no significant block x block order interaction, $F(1,24) = 0.215$, $p = .647$, $\eta^2 = 0.003$, BF_{null} (Bayes Factor in favour of the null hypothesis) = 2.6. Planned t tests on each point shown in Figure 3 (left) showed that the cueing effect was significantly greater than 0 in the congruent block, regardless of block order, t's > 3, p's < .01, but not did not differ from 0 in the incongruent block, t's < 1.65, p's > 1.22. For the unconscious cues, there was crucially no effect of block, $F(1,24) = 0.11$, $p = .743$, $\eta^2 < 0.01$, $BF_{null} = 3.3$, as well as no effect of block order, $F(1,24) = 0.14$, $p = .711$, $\eta^2 < 0.01$, $BF_{null} = 3.4$, and no interaction, $F(1,24) = 0.315$, $p = .58$, $\eta^2 < 0.01$, $BF_{null} = 3.4$.

The results above show that participants adjusted the ways they used the arrow cues in response to the broader context of the task, but only when they were conscious of the cues. However, these analyses were based on average RTs over blocks of 100 trials. It could be that participants do learn from unconsciously presented cues, but that this learning only persists for a short period of time (see Gratton et al., 1992). To test this idea, we grouped trials according by participant, visibility, block, and the congruence of the previous trial, excluding incorrect responses, and trials where the previous response was incorrect, calculated the cueing effect in each group (Figure 5), and subjected this measure to ANOVA analysis. Again, consistent results were obtained conducting cue x previous cue ANOVAs, and using linear mixed models at the trial level.

We first conducted a previous cue x visibility x block x block order mixed ANOVA. This revealed no significant main effect of the previous cue, $F(1, 24) < 0.01$, $p > .95$, $\eta^2 < 0.01$, $BF_{null} = 4.47$, but did show that visibility interacted significantly with the effect of the previous cue, $F(1,24) = 4.382$, $p = .047$, $\eta^2 = 0.011$. Therefore, we once again conducted separate analyses for the conscious and unconscious trials, fitting previous cue x block (mostly congruent or mostly incongruent) x block order (mostly congruent first or mostly incongruent first) mixed ANOVAs to each.

For conscious trials, we found no main effect of previous cue $F(1,24) = 1.504$, $p = .232$, $\eta^2 = 0.01$, $BF_{null} = 2.419$, and no interaction between the previous cue and block $F(1,24) = 1.598$, $p = .218$, $\eta^2 = 0.007$, $BF_{null} = 3.376$. There were however marginally significant interactions previous cue x block order. $F(1,24) = 2.899$, $p = .102$, $\eta^2 = 0.02$, $BF_{null} = 1.545$, and previous cue x block x block order, $F(1,24) = 3.339$, $p = .08$, $\eta^2 = 0.014$, $BF_{null} = 1.616$ (see Figure 3). For unconscious cues, we found no effect of the previous cue $F(1,24) = 1.82$, $p = .19$, $\eta^2 = 0.012$, $BF_{null} = 2.794$, and no interactions of this factor with block order, $F(1,24) = 0.975$, $p = .333$, $\eta^2 = 0.006$, $BF_{null} = 3.056$. Although there was a previous cue x block interaction, $F(1,24) = 5.25$, $p = .031$, $\eta^2 = 0.076$, $BF_{alt} = 8.58$, planned comparisons found no significant effects of the previous cue in either block or for either block order, and the trend was in the opposite direction to that predicted: participants tended to slow down more on trials where the cue pointed in the opposite direction if the previous cue also pointed in the opposite direction (Figure 3). This pattern of results is not consistent with the hypothesis that participants adjusted their beliefs about the cue-target contingencies trial by trial.
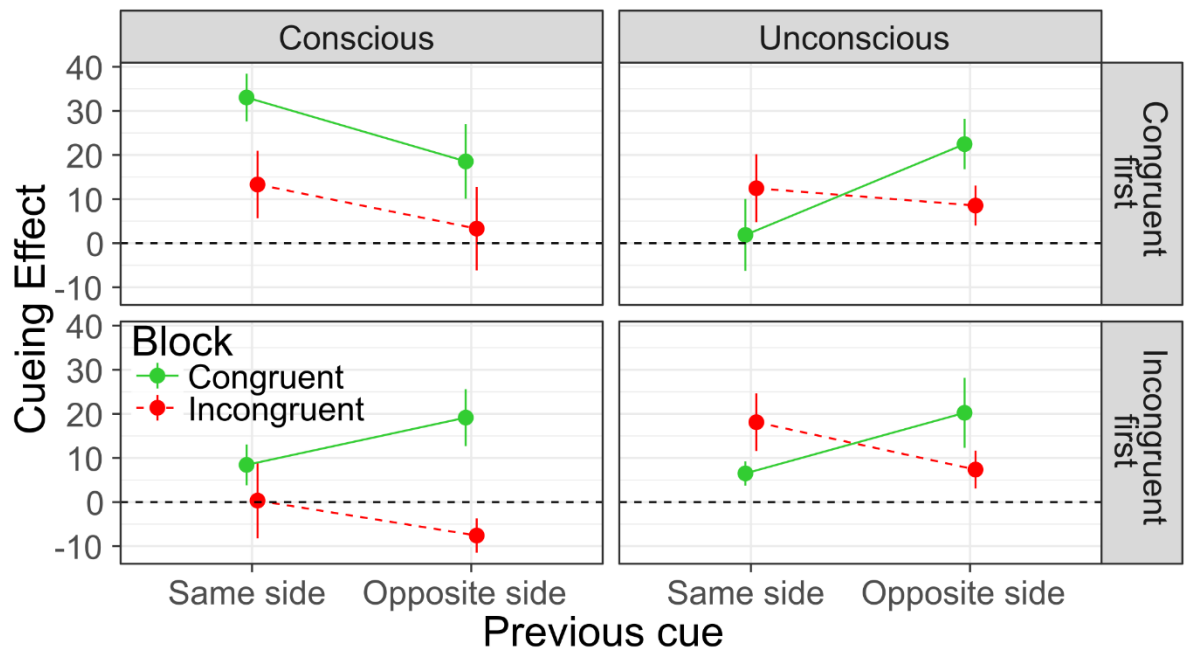
Figure 5. Cueing effect on response times, as a function of the previous cue, cue visibility, block, and block order, Experiment 1. In this and all further figures, error bars show SEM calculated independently for each cell in the design, and so are wider than the actual unexplained variance in our between-subjects analyses.

Visibility Check

Following the exclusion of data from two participants due to poor performance on the visibility check for the conscious cues, reported above, performance was close to ceiling for conscious cues, accuracy = 98%, SD = 11%, d' = 4.38 SD = 1.72, t(25) = 13.493, p < .001, $BF_{alt} > 10^{19}$, and at chance for unconscious cues, accuracy = 50% SD = 11%, d' = -.01 SD = 0.75, t(25) < 0.01, p > .99, BF_{null} = 5.00. Binomial tests on every participant's accuracy scores for the unconsciously presented cues indicated that twenty participants did not differ significantly from chance, two were significantly above chance, and two were significantly below. When asked to indicate to what extent they saw the cues and were not merely guessing, participants gave an average rating of -3.3 for the unconscious cues (SD = 1, max = 0, where -4 = "Completely guessing" and +4 = "Completely saw them"), and 2.9 for the conscious cues (SD = 1.6, min = -2).

There was no relationship between d' sensitivity for unconscious cues in the visibility check and the average cueing effect for each participant, b = -0.12, t(24) = 0.061, p > .95, BF_{null} = 2.753, and a positive intercept, b = 9.39, t(24) = 6.115, p < .001, indicating that subliminal cueing should occur for participants with d' of 0 (indeed, this was the average sensitivity). Importantly, there was also no relationship between d' and the average change in cueing for unconscious cues between congruent

and incongruent blocks, $b = 1.82$, $t(24) = 0.321$, $p > .75$, $BF_{null} = 2.65$, and in this case no significant intercept term, $b = 1.45$, $t(24) = 0.340$, $p > .70$, consistent with our findings above.

Discussion

In line with previous work (e.g. Eimer & Schlaghecken, 2003), we found that left or right arrow cues facilitate their corresponding responses, even when presented subliminally. We also found, consistent with previous work (e.g. Ridderinkhof, 2002; Vossel et al., 2014), that this facilitation effect is greater on blocks where most arrow cues point towards the eventual location of the target A number of previous studies (e.g. Bodner & Dypvik, 2005; Jaśkowski et al., 2003; Klapp, 2007) have also found that this adaptation occurs even when stimuli are presented subliminally. In contrast, we found a significant difference between adaptation to conscious and unconscious stimuli, with Bayesian analyses indicating strong evidence for blockwise learning for conscious cues, but moderate evidence against blockwise learning for unconscious cues. We consider our results in light of previous work in the general discussion below.

Surprisingly, we also found no *Gratton* effects (Gratton et al., 1992): the influence of the cues was not significantly affected by the validity of the previous cue, either in general or for conscious or unconscious cues. However, our complex factorial design (previous cue x block x block order x consciousness) is not ideal for detecting these effects, and so this result should not be interpreted too strongly.

As an aside, we found considerable individual differences in how participants used the consciously presented cues, but not the unconsciously presented ones. This suggest that, when participants were consciously aware of the cues, they explored various strategies for dealing with conflict. These data can be found in the Supplementary Materials.

Experiment 2

In Experiment 1, we investigated a rather extreme form of learning: adjusting to complete reversal of the cue-target contingencies. In Experiment 2, we explore learning of a subtler kind, specifically adjusting beliefs about how predictive the cue is of the target. Additionally, in Experiment 1 participants responded by manually pressing buttons. It could be that learning of some sort does occur for unconsciously presented visual stimuli, but this learning is not broadcast globally, and so does not affect the speed of manual actions. For this reason, in Experiment 2 we monitor participants' eye gaze, and require them to simply saccade to the target as soon as it appears.

We adopt a version of the cuing task reported by Vossel et al. (2014). They presented participants with a series of blocks of 33–38 trials, where in each block the cue pointed to the target location on 50%, 69%, or 88% of trials. They showed that, as the reliability of the cue increased, participants became faster when the cue pointed the right way, and slower when it pointed the wrong way, reflecting learning. They also demonstrated that this learning was best explained by a hierarchical learning model that updates beliefs about both the reliability of the cues and about the degree to which this reliability changed, rather than by a simple reinforcement learning model.

Here, we repeat their experiment, with some changes. First, we replaced their cues with the arrow cues used in Experiment 1. Second, half of our participants were presented with unconscious cues, presented quickly and then masked, while the remainder saw the cues for long enough for them to be clearly visible.

Method

Participants

Twenty-eight participants (mean age = 22.7, SD = 5.3) completed the experiment in a laboratory at the University of London. All reported having normal or corrected-to-normal vision, and received £10 for taking part in the experiment.

Design

Following Vossel et al. (2014) the main phase of the experiment consisted of 12 blocks of 33-38 trials each, for a total of 612 trials. In each block, the cue pointed towards where the target would appear either 50%, 69%, or 88% of the time. There were 4 blocks of each type, and no two consecutive blocks were of the same type. The hierarchical model used by Vossel et al. (2014) is computationally intractable with this number of trials unless the order of trials is held constant across participants (see Vossel et al., 2014, p. 1438) and so for consistency with their work we used the same trial order here.

The main experiment was again followed by a visibility check. Lin and Murray (2014) have suggested that blocked visibility checks of the kind we employed in Experiment 1 can lead participants to stop trying to see the stimuli on the blocks where every trial's stimuli are believed to be unconscious. Therefore, we randomly presented 30 conscious trials and 30 unconscious trials, with the target replaced by a prompt asking "Left or Right?". As in Experiment 1, responses were made using a button box. This check was only completed by the 14 participants who completed the experiment with unconscious cues.

Stimuli & Procedure

Participants were seated in a dimly-lit room, approximately 57 cm from a monitor, with an EyeLink eye-tracker located below the screen. All stimuli were presented in black on a white background, using OpenSesame experiment builder and the PsychoPy back-end, and the PyGaze library was used to communicate with the eye-tracker.

The fixation cross, cues, and mask were unchanged from Experiment 1, except that they now subtended 1.2° vertically and 1.2° (fixation and targets) or 2.4° (cues) horizontally. The cues were offset 2.5° vertically from fixation. Following Vossel et al. (2014), Gabor patches were used as targets, and offset horizontally by 8° from fixation. The possible locations of the targets were marked by empty boxes.

Each trial began with a central fixation, which participants were required to saccade to before proceeding. After fixation was detected, the cross thickened for 250 msec, before being replaced by the cue for 400 (conscious condition) or 49 (unconscious condition) msec, which was then hidden by the mask, which remained onscreen for 34 msec. The target then appeared, and remained onscreen until a saccade away from the fixation was detected. The unconscious cues used here were 16 msec (one frame) longer than in the previous experiment. This was done in part to ensure that the null results obtained in Experiment 1 were not due to the cues being overly-stringently masked, and in part on the basis of pilot data indicating that this presentation time produced robust priming effects without rendering the cues consciously visible.

The procedure for the visibility check was again the same as for the trials, except that participants saw a prompt reading "Left or Right?" instead of the target, and they responded manually.

Eye Movement Data Recording and Analysis

Eye-tracking samples were recorded at 500 Hz, and head movements were compensated for using the EyeLink Remote Camera system. A 9 point calibration sequence was used at the start of the experiment, and after each self-paced break. Our dependant variable, the time taken for participants to initiate a saccade after the onset of the target (RT) was calculated using a custom python script. Saccades were defined as movements over 2º with a velocity in excess of 30º/s, and saccade onsets defined as the point at which pupil velocity exceeded 15% of the maximum velocity achieved during the saccade.

Results

Data Exclusions and Analysis

Participants initiated their eye movements before the onset of the target for 20.4% of conscious cues, and 5.1% of unconscious cues (*pre-emptive responses)*, and furthermore initiated within 100 msec of target onset, and so presumably made their decision before seeing the target, for 12.5% of conscious cues and 2.6% of unconscious (*early responses*). On an additional 2.6% of conscious cues, and 0.6% of unconscious cues, the initial eye movement did not meet our threshold for definition as a proper saccade (*ambiguous responses*). For our analyses of errors we include these trials, but also examine the effect of excluding them on our results. All of these trials were excluded in our analysis of response times, along with a further 1.0% and 1.2% of trials following conscious and unconscious cues respectively with response times greater than 400 msec.

We used multilevel regression models in our analyses in order to treat the proportion of valid cues in a block as a continuous variable, and thus to increase our power to detect learning in the subliminal condition. This approach is also preferable to classical ANOVA in the current context as we were forced to exclude a large number of trials from the response time analysis due to both incorrect responses and cases where participants initiated their eye movements before the onset of the target (see Baayen, Davidon & Bates, 2008). In each analysis, we included a random intercept term for each participant, consistent with classical repeated-measures ANOVA. However, as a number of models failed to converge when additional random effects were included, only participants were treated as a random effect. Furthermore it is not straightforward to calculate Bayes Factors for models with multiple random effects terms.

Errors

Figure 6 (left) shows the proportion of errors on trials with visible and invisible cues, as a function of the cue validity, and the proportion of congruent cues in that block. For conscious cues, but not unconscious cues, participants were more likely to saccade in the wrong direction when the cue pointed in that direction, and this effect increased as the proportion of cues pointing the right way within a block increased. A cue x block ANOVA for the conscious cues indicated a significant main effect of cue direction, $\chi^2 (1) = 983$, $p < .001$, and a cue x block interaction, $\chi^2 (1) = 27.86$, $p < .001$. As the average error rate for the unconscious cues was only 2.9%, it was not appropriate to analyses these trials in the same way. When pre-emptive, early, and ambiguous responses were excluded, the same pattern of results was found, with considerably fewer errors in all conditions (see Supplementary Materials).

Response Times

Figure 6 (right) shows average response times, for visible and invisible cues, as a function of the cue validity, and the proportion of congruent cues that block, after excluding error trials and the other problematic responses, discussed above. An initial cue x block x visibility model revealed a significant interaction between cue and visibility, $\chi^2(1) = 121.6$, $p < .001$, indicating that the cueing effect was stronger for conscious cues, a significant interaction between cue and block, $\chi^2 (1) = 7.6$, $p = .006$, indicating greater cueing effects on blocks with more congruent cues, and, crucially, a cue x block x visibility interaction, $\chi^2 (1) = 6.8$, $p = .009$. To unpick this, we again fit separate cue x block models for the conscious and unconscious cues. For conscious cues, we found significant effects of cue, $\chi^2 (1) = 256.3$, $p < .001$, of block, and, importantly, their interaction, $\chi^2 (1) = 10.4$, $p < .001$, $BF_{alt} = 8.9$: participants were faster for congruent cues, faster on blocks with more congruent cues, and were more influenced by the cues on blocks where more of the cues were congruent. For unconscious cues, we found a significant effect of cue, $\chi^2 (1) = 39.9$, $p < .001$. The crucial cue x block interaction

was not significant, $\chi^2(1) = 0.3$, $p = .571$, $BF_{null} = 21.6$. Consistent results were found when also including actual block number (1 to 12) as a linear predictor to allow for changes in cue use over time, and the crucial cue x block interaction for unconscious cues remained non-significant, $\chi^2(1) = 0.5$, $p = .464$, $BF_{null} = 14.4$.

As in Experiment 1, we wished to investigate if learning could occur for unconscious cues across shorter time-scales, specifically, between one trial and the next. We again coded the data according to the congruence of the cue on the previous trial, excluding incorrect responses, trials where the previous response was incorrect, and the other problematic trials discussed above. Figure 7 shows the effect of the current cue and the previous cue on both errors and response times, for conscious and unconscious cues. As expected for conscious cues there was a significant cue x previous cue interaction, $\chi^2(1) = 25.0$, $p < .001$, as the effect of the current cue was strongest when the previous cue was congruent. Again, there were insufficient errors for unconscious cues for a proper analysis, but the mean error rates were close to identical. For response times, we found a significant cue x previous cue x visibility interaction, $\chi^2(1) = 8.8$, $p = .003$. Follow-up models revealed a significant cue x previous cue interaction, in the expected direction, for conscious cues, $\chi^2(1) = 14.0$, $p < .001$, $BF_{alt} = 32.4$, but, crucially, not for unconscious cues, $\chi^2(1) = 0.5$, $p = .504$, $BF_{null} = 20.6$.

Figure 6. Errors (left) and response times for correct responses (right) for conscious and unconscious cues in Experiment 2, as functions of the congruence of the cue, and the proportion of congruent cues within a given block.
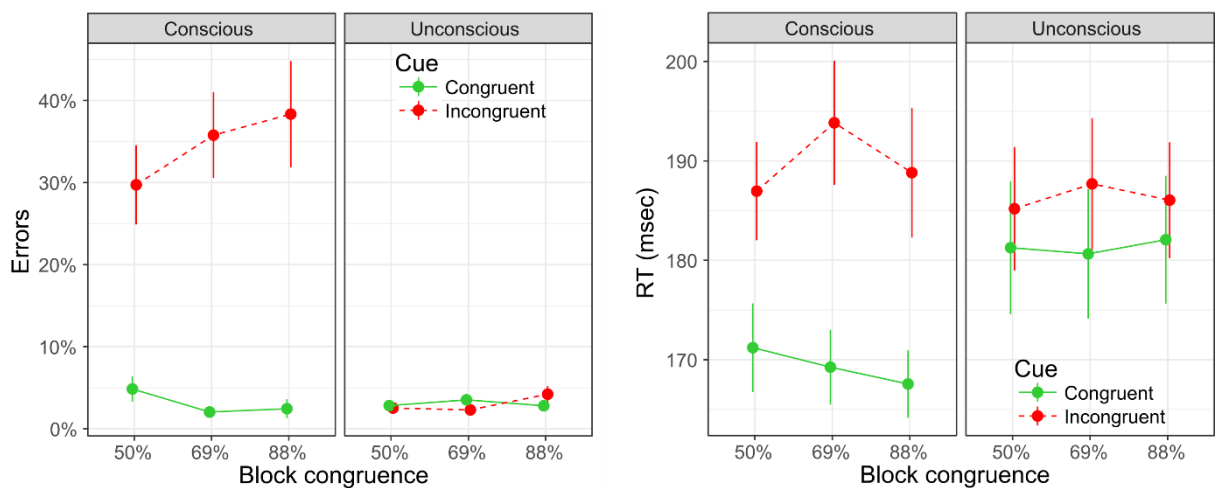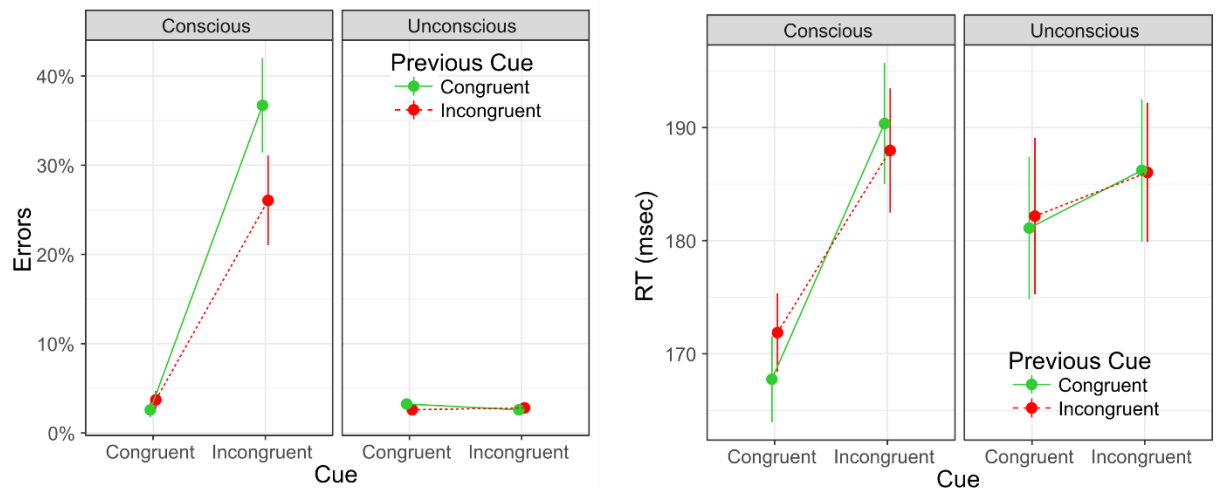
Figure 7. Errors (left) and response times for correct responses (right) for conscious and unconscious cues in Experiment 2, as functions of the congruence of the cue, and congruence of the cue on the previous trial.

Visibility check

Performance was near ceiling in the visibility check for the 400 msec cues, accuracy = 96% (20%), d' = 4.0 (1.1), t(13) = 14.077, p < .001. Performance for the 50 msec cues was significantly above chance, accuracy = 65% (14%), d' = 0.9 (0.8), t(13) = 4.266, p = .001. Binomial tests on every participants' accuracy score were significantly above chance for conscious cues for all participants, and above chance for 7 of 14 participants for unconscious cues. Therefore participants could to some degree discriminate between leftward and rightward cues that were intended to be unconscious.

To estimate the unconscious cueing effect for each participant, we subtracted their response times for congruent cues from their times for congruent cues, as in Experiment 1. This effect was not predicted by participants' d' score for unconscious cues in the visibility test, b = -0.5, t(12) = .860, BF$_{null}$ = 2.21, and was in fact slightly negatively related to sensitivity on the visibility test. As an estimate of the influence of block congruence on the subliminal cueing effect, we calculated cueing effects in the same way separately for each block, and subtracted the effect for the 88% congruent block from that for the 50% block. This measure was also not predicted by d' scores, b = -1.317, t(12) = 0.352, p = .741, BF$_{null}$ = 2.15, and was again actually negatively correlated with sensitivity. Therefore, although participants performed above chance on the visibility test for supposedly unconscious cues, individual difference in performance on this test predicted neither simple cueing effects in the main experiment nor the modulation of these effects by learning. This suggests that the apparently visibility of the cues in the post test did not affect their use in the main experiment.

Discussion

The current results broadly replicate those of Experiment 1. Participants were faster to respond on trials where the cue was congruent with the target, for both unconscious and conscious cues. For conscious cues only, there was also evidence of learning between blocks, as participants relied more on the cues during blocks where a greater proportion of the cues were congruent. In contrast to Experiment 1, but consistent with previous work, similar results were obtained for learning from trial to trial, as participants relied more on the cue when the previous cue was congruent for conscious cues, but not for unconscious cues.

One potential issue for this experiment is that performance on the visibility check for supposedly unconscious cues was above chance overall, and above chance for 7 of the 14 participants. Therefore, participants may have had access to some information about the cues on some trials, despite the absence of learning when cues were masked. This makes our experiments a conservative test of the absence of unconscious learning. We will return to this issue in the General Discussion.

General Discussion

Across two experiments, we have replicated the classic finding that attention can be directed by both conscious and unconscious cues (e.g. Eimer & Schlaghecken, 1998; Jiang et al., 2013; Schlaghecken & Eimer, 2002). However, while participants showed signs of learning rapidly about cues that they were consciously aware of, changing strategy when the cue-target contingencies reversed in Experiment 1, or adjusting the credence they gave the cues in Experiment 2, we found no evidence of learning about cues that were processed unconsciously. Furthermore, these differences between the conscious and unconscious conditions were significant. Similarly, while the way participants used a cue on a given trial was affected by the accuracy of the cue on the previous trial when the cues were conscious, this was not consistently the case when the cues were presented unconsciously. Taken together, these results suggest that while some information can be processed and affect behaviour without reaching conscious awareness, effective rapid learning, at least on this particular task, within the time frames we investigated of up to 100 trials, requires that information is represented consciously. That is, consciousness facilitated learning.

Some studies of conflict adaptation have reported blockwise learning effects for subliminal cues (Bodner & Dypvik, 2005; Bodner et al., 2006; Jaśkowski et al., 2003, Klapp, 2007). However, the tasks used in these studies were very different from ours (parity judgement, lexical decision, gap detection) and the statistical analyses and tests for visibility were not always ideal. The study of this type that most closely resembles our design, using different cue contingencies within participants (Bodner and Stalinski, 2008), found no unconscious learning for participants in a condition equivalent

to ours with no cognitive load. Why, then, might blockwise learning be found for unconscious cues in some studies, but not others? One possibility is that control of the conscious visibility of the stimuli in previous studies was inadequate (see Vadillo et al., 2015). In particular, while some studies have shown that individual differences in conscious visibility of the stimuli did not predict simple priming, and so that evidence for priming is not due to a subset of participants who are aware of the stimuli, to our knowledge none have tested the relationship between visibility and learning in the same way. A more interesting possibility is that the key difference is that between a between-participant design where all participants see either high- or low-validity cues, or a within-participants design where cue validity changes over the course of the experiment. It seems to us that the latter can be better described as capturing *learning* than the former can, but it is in the within-participant experiments that evidence of learning has not been consistently found. Having said this, it is not clear what mechanism might underlie the differences between conditions in between-participant designs, but the evidence available suggests that they may not be due to learning or adaptation.

A number of other studies have also found that the validity of the previous cue does not alter the processing of the cue on the current trial when cues are presented unconsciously (Frings & Wentura, 2008; Kunde, 2003; but see also van Gaal, Lamme, & Ridderinkhof, 2010 for a conflicting result), in line with the results of Experiment 2. As noted above, this effect was not found for conscious or unconscious cues in Experiment 1, but this study was not ideally-designed to find such an effect. We would also note that it is not totally clear how unconscious learning across an entire experimental block might take place, if not as the accumulation of minor adjustments made on a trial-by-trial basis.

We would also note that our results do not go directly against more recent work on context-specific conflict adaptation (e.g. Panadero et al., 2015; Reuss et al., 2014; Schouppe et al., 2014). In these studies participants can benefit from a conscious or unconscious cue indicating the context (high or low conflict) in which they are working, but must infer which context is which from the themselves. , There were no such signals about conflict in our study, so it remains possible that participants may be able to unconsciously adjust to the prevalence of conflict in different contexts when this is signalled, but do not do so within a block if there is no such signal. Our results do not directly address this question, and so we will not discuss it further.

As noted above, performance on the visibility checks for Experiment 2 did not provide strong evidence that participants were not consciously aware of the arrow cues intended to be unconscious. However, we argue that this limitation does not invalidate our results for two reasons. First, the visibility checks do not demonstrate that participants were consciously aware of the cues either: a number of studies have shown that above chance discrimination can occur in the absence of conscious awareness of the stimuli presented (Forster & Govier, 1978; Kunimoto, Miller, & Pashler, 2001;

Lamy, Salti, & Bar-Haim, 2009; Persaud & McLeod, 2008). Second, our aim was to investigate whether learning occurs in this kind of task when the cues are unconscious. The possibility that some subjects were conscious of some cues in some trials makes it more, not less likely that learning would occur, so our experiments are a conservative test of the absence of unconscious learning. Our data consistently show that participants do learn about the conscious cues, but do not learn about the masked cues. That finding supports the conclusion that learning did not occur when the cues were not consciously represented. This pattern of results would not be expected if both kinds of cue were actually processed consciously. More generally, our visibility check results would be problematic if we were to make claims about what can be done with unconscious representations, rather than claims about what cannot. We take it as known that unconscious cues can prime behaviour (Boy & Sumner, 2010; Schlaghecken & Eimer, 2000; Sumner et al., 2007; Wenke et al., 2010).

We would stress that these results should not be over interpreted. As noted above, it is not clear that consciousness is strictly necessary for any task. It may be that our participants learned both consciously and unconsciously, but that unconscious learning here was ineffective, and required either more experience than the 100 trials per block we allowed participants in Experiment 1, or a more sensitive behavioural measure. Similarly, we do not claim, on the basis of our data, that conscious awareness of stimuli is necessary for all meaningful kinds of learning; this experiment focused on a very restricted form of learning, about the best ways for participants to shift their attention in response to visual cues on our task. Indeed, it should be clear, *a priori*, that consciousness cannot be necessary for learning more generally: no one would argue that the evolution of conscious awareness occurred before that of Hebbian learning. Rather, our claim here is that conscious and unconscious information processing differ in terms of how they facilitate learning. Participants showed no evidence of learning from unconscious cues, but were able to learn from a single exposure to a conscious cue, indicating that whatever the nature of unconscious processing here, it differs significantly from conscious processing.

Another claim we do not wish to make is that these results show that learning must be a deliberate process, rather than an automatic one. Shea and Frith (2016) argued that the distinction between conscious and unconscious representations is orthogonal to the distinction between deliberate and automatic processes. Automatic process can apply to conscious or unconscious representations, and it remains unclear whether deliberate processes are restricted so that they can only operate on conscious representations. From this perspective, our data indicate that learning on our task occurred only for conscious representations, but say little about whether the learning found in the conscious case is automatic or deliberate.

With that said, these results clearly have implications for theories of consciousness. First, they are inconsistent with accounts that hold that consciousness is epiphenomenal (Baumeister et al., 2011;

Huxley, 1874; Pockett, 2004), or, in other words, that the only difference between conscious and unconscious information processing is that we are subjectively aware of the former but not the latter. Instead, they indicate qualitative differences between conscious and unconscious processing of the same information. In our specific case at least, these results indicate that learning is facilitated by conscious awareness of the stimuli.

According to one theory of consciousness (Dehaene & Changeux, 2011; Salti et al., 2015), conscious and unconscious processing both involve similar processes over a range of cortical systems, but differ in terms of their neural dynamics. Unconscious processes are characterised by transient, largely feed-foward passages of information through cortical systems. Beyond a critical point, however, non-linear recurrent loops become engaged, allowing information to be maintained and updated over longer periods of time, and to be broadcast to disparate cortical systems via a global workspace. On our task, from this point of view, the initial passage of information about the cues through the visual system is sufficient to guide participants' attention. However, the transition to conscious, recurrent processing allows for the rapid adjustment of higher-order beliefs, or learning. This could be the case for two reasons. It may be that feedforward activation is simply too short lived for higher-order learning about cue-target contingencies to occur (see Draine & Greenwald, 1998). Alternatively, it may be global broadcast of information is necessary for effective learning here — information about unconsciously processed cues may not be accessible beyond the visuomotor system, whereas consciously processed cues are accessible to higher-order systems necessary for this kind of learning.

Finally, it is worth considering that the nature of our task likely renders a whole swathe of processing unconscious. When presented with unconscious cues, participants a) process these cues without awareness, b) likely predict the target location and reallocate their attention without being aware of doing so, and c) while they are conscious of the target location, they are not aware if this confirms or violates their unconscious prediction. By contrast, Pessiglione et al. (2008) presented participants with masked cues, and rewarded or punished their subsequent overt responses (button presses). They found that participants did learn the appropriate cue-response contingencies, indicating that learning about unconsciously-presented cues can occur when participants are conscious of both their responses and their consequences. Therefore, a task for future research will be to dissociate these components of the learning process to discover what information it is necessary or sufficient that be participants be conscious of for learning to occur: the cue, their response to the cue, or feedback about the appropriateness of their response.

References

Ansorge, U., Kunde, W., & Kiefer, M. (2014). Unconscious vision and executive control: How unconscious processing and conscious action control interact. *Consciousness and Cognition*, *27*, 268–287. https://doi.org/10.1016/j.concog.2014.05.009

Atas, A., Faivre, N., Timmermans, B., Cleeremans, A., & Kouider, S. (2014). Nonconscious Learning From Crowded Sequences. *Psychological Science*, *25*(1), 113–119. https://doi.org/10.1177/0956797613499591

Baumeister, R. F., Masicampo, E. J., & Vohs, K. D. (2011). Do Conscious Thoughts Cause Behavior? *Annual Review of Psychology*, *62*(1), 331–361. https://doi.org/10.1146/annurev.psych.093008.131126

Bekinschtein, T. A., Dehaene, S., Rohaut, B., Tadel, F., Cohen, L., & Naccache, L. (2009). Neural signature of the conscious processing of auditory regularities. *Proceedings of the National Academy of Sciences*, *106*(5), 1672–1677. https://doi.org/10.1073/pnas.0809667106

Blais, C., Harris, M. B., Guerrero, J. V., & Bunge, S. A. (2012). Rethinking the role of automaticity in cognitive control. *The Quarterly Journal of Experimental Psychology*, *65*(2), 268–276. https://doi.org/10.1080/17470211003775234

Bodner, G. E., & Dypvik, A. T. (2005). Masked priming of number judgments depends on prime validity and task. *Memory & Cognition*, *33*(1), 29–47.

Bodner, G. E., & Masson, M. E. (2003). Beyond spreading activation: An influence of relatedness proportion on masked semantic priming. *Psychonomic Bulletin & Review*, *10*(3), 645–652.

Bodner, G. E., & Masson, M. E. (2004). Beyond binary judgments: Prime validity modulates masked repetition priming in the naming task. *Memory & Cognition*, *32*(1), 1–11.

Bodner, G. E., & Masson, M. E. J. (2001). Prime Validity Affects Masked Repetition Priming: Evidence for an Episodic Resource Account of Priming. *Journal of Memory and Language*, *45*(4), 616–647. https://doi.org/10.1006/jmla.2001.2791

Bodner, G. E., Masson, M. E. J., & Richard, N. T. (2006). Repetition proportion biases masked priming of lexical decisions. *Memory & Cognition*, *34*(6), 1298–1311. https://doi.org/10.3758/BF03193273

Bodner, G. E., & Mulji, R. (2010). Prime proportion affects masked priming of fixed and free-choice responses. *Experimental Psychology*, *57*(5), 360–366. https://doi.org/10.1027/1618-3169/a000043

27

Bodner, G. E., & Stalinski, S. M. (2008). Masked repetition priming and proportion effects under cognitive load. *Canadian Journal of Experimental Psychology/Revue Canadienne de Psychologie Expérimentale*, *62*(2), 127.

Boy, F., & Sumner, P. (2010). Tight coupling between positive and reversed priming in the masked prime paradigm. *Journal of Experimental Psychology. Human Perception and Performance*, *36*(4), 892–905. https://doi.org/10.1037/a0017173

Byrne, R. M. J. (2016). Counterfactual Thought. *Annual Review of Psychology*, *67*(1), 135–157. https://doi.org/10.1146/annurev-psych-122414-033249

Clark, R. E. (1998). Classical Conditioning and Brain Systems: The Role of Awareness. *Science*, *280*(5360), 77–81. https://doi.org/10.1126/science.280.5360.77

Crump, M. J., Gong, Z., & Milliken, B. (2006). The context-specific proportion congruent Stroop effect: Location as a contextual cue. *Psychonomic Bulletin & Review*, *13*(2), 316–321.

Dehaene, S., & Changeux, J.-P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, *70*(2), 200–227. https://doi.org/10.1016/j.neuron.2011.03.018

Dehaene, S., Charles, L., King, J.-R., & Marti, S. (2014). Toward a computational theory of conscious processing. *Current Opinion in Neurobiology*, *25*, 76–84. https://doi.org/10.1016/j.conb.2013.12.005

Dehaene, S., Kerszberg, M., & Changeux, J.-P. (1998). A neuronal model of a global workspace in effortful cognitive tasks. *Proceedings of the National Academy of Sciences*, *95*(24), 14529–14534. https://doi.org/10.1073/pnas.95.24.14529

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*(1–2), 1–37. https://doi.org/10.1016/S0010-0277(00)00123-2

Desender, K., & Van Den Bussche, E. (2012). Is Consciousness Necessary for Conflict Adaptation? A State of the Art. *Frontiers in Human Neuroscience*, *6*. https://doi.org/10.3389/fnhum.2012.00003

Draine, S. C., & Greenwald, A. G. (1998). Replicable unconscious semantic priming. *Journal of Experimental Psychology: General*, *127*(3), 286. https://doi.org/10.1037/0096-3445.127.3.286

Eimer, M., & Schlaghecken, F. (1998). Effects of masked stimuli on motor activation: behavioral and electrophysiological evidence. *Journal of Experimental Psychology: Human Perception and Performance*, *24*(6), 1737.

Eimer, M., & Schlaghecken, F. (2003). Response facilitation and inhibition in subliminal priming. *Biological Psychology*, *64*(1–2), 7–26. https://doi.org/10.1016/S0301-0511(03)00100-5

Eriksen, B. A., & Eriksen, C. W. (1974). Effects of noise letters upon the identification of a target letter in a nonsearch task. *Perception & Psychophysics*, *16*(1), 143–149. https://doi.org/10.3758/BF03203267

Fanselow, M. S., & Poulos, A. M. (2005). The Neuroscience of Mammalian Associative Learning. *Annual Review of Psychology*, *56*(1), 207–234. https://doi.org/10.1146/annurev.psych.56.091103.070213

Forster, P. M., & Govier, E. (1978). Discrimination without awareness? *Quarterly Journal of Experimental Psychology*, *30*(2), 289–295. https://doi.org/10.1080/14640747808400676

Frings, C., & Wentura, D. (2008). Trial-by-trial effects in the affective priming paradigm. *Acta Psychologica*, *128*(2), 318–323. https://doi.org/10.1016/j.actpsy.2008.03.004

Frith, C. (2010). What is consciousness for? *Pragmatics & Cognition*, *18*(3), 497–551.

Frith, C. D., & Metzinger, T. K. (2016). What's the use of Consciousness? How the Stab of Conscience Made Us Really Conscious. In A. K. Engel, K. J. Friston, & D. Kagic (Eds.), *The Pragmatic Turn: Toward Action-Oriented Views in Cognitive Science*. Cambridge, MA, USA: MIT Press. Retrieved from http://www.blogs.uni-mainz.de/fb05philosophie/files/2013/04/Frith_Metzinger_Regret_2016_penultimate.pdf

Gagliano, M., Vyazovskiy, V. V., Borbély, A. A., Grimonprez, M., & Depczynski, M. (2016). Learning by Association in Plants. *Scientific Reports*, *6*, 38427. https://doi.org/10.1038/srep38427

Gratton, G., Coles, M. G., & Donchin, E. (1992). Optimizing the use of information: strategic control of activation of responses. *Journal of Experimental Psychology. General*, *121*(4), 480–506.

Heinemann, A., Kunde, W., & Kiesel, A. (2009). Context-specific prime-congruency effects: On the role of conscious stimulus representations for cognitive control. *Consciousness and Cognition*, *18*(4), 966–976. https://doi.org/10.1016/j.concog.2009.08.009

Huxley, T. H. (1874). On the hypothesis that animals are automata, and its history. *Nature*, *10*, 362–366.

Jacoby, L. L., Lindsay, D. S., & Hessels, S. (2003). Item-specific control of automatic processes: stroop process dissociations. *Psychonomic Bulletin & Review*, *10*(3), 638–644.

Jaśkowski, P., Skalska, B., & Verleger, R. (2003). How the self controls its 'automatic pilot' when processing subliminal information. *Journal of Cognitive Neuroscience*, *15*(6), 911–920.

Jiang, J., van Gaal, S., Bailey, K., Chen, A., & Zhang, Q. (2013). Electrophysiological correlates of block-wise strategic adaptations to consciously and unconsciously triggered conflict. *Neuropsychologia*, *51*(13), 2791–2798. https://doi.org/10.1016/j.neuropsychologia.2013.09.020

Kiss, M., & Eimer, M. (2008). ERPs reveal subliminal processing of fearful faces. *Psychophysiology*, *45*(2), 318–326. https://doi.org/10.1111/j.1469-8986.2007.00634.x

Klapp, S. T. (2007). Nonconscious control mimics a purposeful strategy: Strength of Stroop-like interference is automatically modulated by proportion of compatible trials. *Journal of Experimental Psychology: Human Perception and Performance*, *33*(6), 1366. https://doi.org/10.1037/0096-1523.33.6.1366

Kouider, S., & Dehaene, S. (2007). Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *362*(1481), 857–875. https://doi.org/10.1098/rstb.2007.2093

Kunde, W. (2003). Sequential modulations of stimulus-response correspondence effects depend on awareness of response conflict. *Psychonomic Bulletin & Review*, *10*(1), 198–205. https://doi.org/10.3758/BF03196485

Kunimoto, C., Miller, J., & Pashler, H. (2001). Confidence and Accuracy of Near-Threshold Discrimination Responses. *Consciousness and Cognition*, *10*(3), 294–340. https://doi.org/10.1006/ccog.2000.0494

Lamy, D., Salti, M., & Bar-Haim, Y. (2009). Neural correlates of subjective awareness and unconscious processing: an ERP study. *Journal of Cognitive Neuroscience*, *21*(7), 1435–1446.

Lovibond, P. F., & Shanks, D. R. (2002). The role of awareness in Pavlovian conditioning: Empirical evidence and theoretical implications. *Journal of Experimental Psychology: Animal Behavior Processes*, *28*(1), 3–26. https://doi.org/10.1037//0097-7403.28.1.3

Mathys, C., Daunizeau, J., Friston, K. J., & Stephan, K. E. (2011). A Bayesian Foundation for Individual Learning Under Uncertainty. *Frontiers in Human Neuroscience*, *5*. https://doi.org/10.3389/fnhum.2011.00039

McClelland, J. L., Rumelhart, D. E., & PDP Research Group, C. (Eds.). (1986). *Parallel Distributed Processing: Explorations in the Microstructure, Vol. 1: Foundations*. Cambridge, MA, USA: MIT Press.

McCormick, P. A. (1997). Orienting attention without awareness. *Journal of Experimental Psychology: Human Perception and Performance*, *23*(1), 168–180. https://doi.org/10.1037/0096-1523.23.1.168

Norman, D. A., & Shallice, T. (1986). Attention to Action. In R. J. Davidson, G. E. Schwartz, & D. Shapiro (Eds.), *Consciousness and Self-Regulation* (pp. 1–18). Springer US. https://doi.org/10.1007/978-1-4757-0629-1_1

Panadero, A., Castellanos, M. C., & Tudela, P. (2015). Unconscious context-specific proportion congruency effect in a stroop-like task. *Consciousness and Cognition*, *31*, 35–45. https://doi.org/10.1016/j.concog.2014.09.016

Persaud, N., & McLeod, P. (2008). Wagering demonstrates subconscious processing in a binary exclusion task. *Consciousness and Cognition*, *17*(3), 565–575. https://doi.org/10.1016/j.concog.2007.05.003

Pessiglione, M., Petrovic, P., Daunizeau, J., Palminteri, S., Dolan, R. J., & Frith, C. D. (2008). Subliminal Instrumental Conditioning Demonstrated in the Human Brain. *Neuron*, *59*(4), 561–567. https://doi.org/10.1016/j.neuron.2008.07.005

Pockett, S., Banks, W. P., & Gallagher, S. (Eds.). (2009). *Does Consciousness Cause Behavior?* Cambridge, Mass: The MIT Press.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, *32*(1), 3–25. https://doi.org/10.1080/00335558008248231

Reuss, H., Desender, K., Kiesel, A., & Kunde, W. (2014). Unconscious conflicts in unconscious contexts: The role of awareness and timing in flexible conflict adaptation. *Journal of Experimental Psychology: General*, *143*(4), 1701–1718. https://doi.org/10.1037/a0036437

Reuss, H., Pohl, C., Kiesel, A., & Kunde, W. (2011). Follow the sign! Top-down contingent attentional capture of masked arrow cues. *Advances in Cognitive Psychology*, *7*(1), 82–91. https://doi.org/10.2478/v10053-008-0091-3

Salti, M., Monto, S., Charles, L., King, J.-R., Parkkonen, L., & Dehaene, S. (2015). Distinct cortical codes and temporal dynamics for conscious and unconscious percepts. *eLife*, e05652. https://doi.org/10.7554/eLife.05652

Schlaghecken, F., & Eimer, M. (2000). A central-peripheral asymmetry in masked priming. *Perception & Psychophysics*, *62*(7), 1367–1382. https://doi.org/10.3758/BF03212139

Schouppe, N., de Ferrerre, E., Van Opstal, F., Braem, S., & Notebaert, W. (2014). Conscious and unconscious context-specific cognitive control. *Frontiers in Psychology*, *5*. https://doi.org/10.3389/fpsyg.2014.00539

Seitz, A. R., & Watanabe, T. (2003). Psychophysics: Is subliminal learning really passive? *Nature*, *422*(6927), 36–36. https://doi.org/10.1038/422036a

Shea, N., & Frith, C. D. (2016). Dual-process theories and consciousness: the case for 'Type Zero' cognition: Table 1. *Neuroscience of Consciousness*, *2016*(1), niw005. https://doi.org/10.1093/nc/niw005

Simon, J. R., & Wolf, J. D. (1963). Choice reaction time as a function of angular stimulus-response correspondence and age. *Ergonomics*, *6*(1), 99–105.

Sokolov, E. N. (1963). Higher Nervous Functions: The Orienting Reflex. *Annual Review of Physiology*, *25*(1), 545–580. https://doi.org/10.1146/annurev.ph.25.030163.002553

Stroop, J. R. (1935). Studies of interference in serial verbal reactions. *Journal of Experimental Psychology*, *18*(6), 643–662. https://doi.org/10.1037/h0054651

Sumner, P., Nachev, P., Morris, P., Peters, A. M., Jackson, S. R., Kennard, C., & Husain, M. (2007). Human Medial Frontal Cortex Mediates Unconscious Inhibition of Voluntary Action. *Neuron*, *54*(5), 697–711. https://doi.org/10.1016/j.neuron.2007.05.016

Tenenbaum, J. B., Kemp, C., Griffiths, T. L., & Goodman, N. D. (2011). How to grow a mind: Statistics, structure, and abstraction. *Science*, *331*(6022), 1279–1285.

Tononi, G. (2004). An information integration theory of consciousness. *BMC Neuroscience*, *5*, 42. https://doi.org/10.1186/1471-2202-5-42

Vadillo, M. A., Konstantinidis, E., & Shanks, D. R. (2015). Underpowered samples, false negatives, and unconscious learning. *Psychonomic Bulletin & Review*, *23*(1), 87–102. https://doi.org/10.3758/s13423-015-0892-6

van Gaal, S., Lamme, V. A. F., & Ridderinkhof, K. R. (2010). Unconsciously Triggered Conflict Adaptation. *PLoS ONE*, *5*(7), e11508. https://doi.org/10.1371/journal.pone.0011508

van Gaal, S., Naccache, L., Meuwese, J. D., van Loon, A. M., Leighton, A. H., Cohen, L., & Dehaene, S. (2014). Can the meaning of multiple words be integrated unconsciously? *Philosophical Transactions of the Royal Society of London B: Biological Sciences*, *369*(1641), 20130212.

van Gaal, S., Ridderinkhof, K. R., van den Wildenberg, W. P. M., & Lamme, V. A. F. (2009). Dissociating consciousness from inhibitory control: evidence for unconsciously triggered response inhibition in the stop-signal task. *Journal of Experimental Psychology. Human Perception and Performance*, *35*(4), 1129–1139. https://doi.org/10.1037/a0013551

Vossel, S., Mathys, C., Daunizeau, J., Bauer, M., Driver, J., Friston, K. J., & Stephan, K. E. (2014). Spatial Attention, Precision, and Bayesian Inference: A Study of Saccadic Response Speed. *Cerebral Cortex*, *24*(6), 1436–1450. https://doi.org/10.1093/cercor/bhs418

Webb, T. W., Kean, H. H., & Graziano, M. S. A. (2016). Effects of Awareness on the Control of Attention. *Journal of Cognitive Neuroscience*, 1–10. https://doi.org/10.1162/jocn_a_00931

Wenke, D., Fleming, S. M., & Haggard, P. (2010). Subliminal priming of actions influences sense of control over effects of action. *Cognition*, *115*(1), 26–38. https://doi.org/10.1016/j.cognition.2009.10.016

Winkielman, P., & Berridge, K. C. (2004). Unconscious Emotion. *Current Directions in Psychological Science*, *13*(3), 120–123. https://doi.org/10.1111/j.0963-7214.2004.00288.x

Wokke, M. E., van Gaal, S., Scholte, H. S., Ridderinkhof, K. R., & Lamme, V. A. F. (2011). The Flexible Nature of Unconscious Cognition. *PLoS ONE*, *6*(9), e25729. https://doi.org/10.1371/journal.pone.0025729