

A new empirical challenge for local theories of consciousness

Matthias Michel^{1,2} & Adrien Doerig³

1. Consciousness, Cognition & Computation Group, Université Libre de Bruxelles, Belgium
2. Mind, Brain and Consciousness Center, New York University, USA
3. Laboratory of Psychophysics, Brain Mind Institute, EPFL, Switzerland

Authors' version. Forthcoming in *Mind & Language*. Please cite the published version.

Abstract: Local theories of consciousness state that one is conscious of a feature if it is adequately represented and processed in sensory brain areas, given some background conditions. We challenge the core prediction of local theories based on recently discovered long-lasting postdictive effects demonstrating that features can be represented for hundreds of milliseconds in perceptual areas without being consciously perceived. Unlike previous empirical data aimed against local theories, proponents of local theories cannot explain these effects away by conjecturing that subjects are phenomenally conscious of features that they cannot report. Only a strong and counterintuitive version of this claim can account for long-lasting postdictive effects. Although possible, we argue that adopting this strong version of the “overflow hypothesis” would have the effect of nullifying the weight of the evidence taken to support local theories of consciousness in the first place. We also discuss several alternative explanations that proponents of local theories could offer.

Introduction

The debate between local and global theories of consciousness is all the rage at the moment. Localists hold that, given some background conditions, neural activity within sensory modules can give rise to conscious experiences (Block, 2007; Lamme, 2006, 2010, 2015; Zeki & Bartels, 1999; Zeki, 2003). For instance, according to the local recurrence theory, reentrant activity within the visual system is necessary and sufficient for conscious visual experiences (Lamme, 2006)¹. Globalists defend that consciousness involves the large-scale coordination of a variety of neuro-cognitive modules (Dehaene & Changeux, 2011), or a set of high-level cognitive functions such as the capacity to form higher-order thoughts about one's perceptual states (Brown et al., 2019; Lau & Rosenthal, 2011).

One's view on this matter has far-reaching theoretical consequences. Localists tend to believe that consciousness is rich (Bronfman et al., 2014; Sligte et al., 2010; Lamme, 2010; Landman, 2003), that it does not require attention (Koch & Tsuchiya, 2003; Lamme, 2003, 2004; Pinto et al., 2017), and that phenomenal consciousness overflows cognitive access (Block, 2007, 2011a; Lamme, 2010). Globalists typically hold that consciousness is sparse (Cohen et al., 2014, 2016; Kouider et al., 2010; Ward et al., 2016), requires attention (Cohen et al., 2012; Mack & Rock, 1998; Mack & Clarke, 2012), and is co-extensive with cognitive access (Cohen & Dennett, 2011; Dehaene et al., 2006; Naccache, 2018).

So far the debate has focused on two main issues: determining whether phenomenal consciousness overflows cognitive access (Block, 2007, 2011a; Cohen et al., 2016; Knotts et al., 2019; Odegaard et al., 2018), and whether the prefrontal cortex is involved in consciousness of contents (Michel & Morales, 2019; Odegaard et al., 2017; Pitts et al., 2014a, 2014b; Tsuchiya et al., 2015). In this article we put forward a new phenomenon that should be accounted for by current theories of consciousness: long-lasting postdiction. In long-lasting postdiction, one's conscious perception of a stimulus depends on events happening more than 300 milliseconds *after* stimulus presentation.

After presenting local theories in Section 1, we discuss long-lasting postdiction in Section 2. In Section 3, we surmise that long-lasting postdiction constitutes a new empirical challenge for local theories of consciousness. Finally, we consider a possible localist response: subjects could have *phenomenally conscious* but *unaccessed* experiences. We argue that adopting this response would commit localists to a strong, and counter-intuitive interpretation of the overflow hypothesis. Moreover, adopting this interpretation would nullify the weight of the evidence interpreted as supporting local theories in the first place.

¹ We focus mainly on conscious *visual* perception in this article.

1. Local Theories of Consciousness

According to local views, a perceptual feature is consciously experienced when it is appropriately represented in sensory systems, given some background conditions². As localism is a broad family of theories, what “appropriately” means depends on the local theory under consideration. Here, we consider only two of the most popular local theories: the micro-consciousness theory (Zeki & Bartels, 1999; Zeki, 2003), and the local recurrence theory (Lamme, 2006, 2010, 2015), focusing on the latter.

According to the micro-consciousness theory “processing sites are also perceptual sites” (Zeki, 2003; p.214). This theory is *extremely* local. The simple fact of representing a perceptual feature is sufficient for being conscious of that feature, given some background conditions. One becomes conscious of individual visual features before integrating them into a coherent whole.

According to the local recurrence theory, consciousness depends on "recurrent" activity (i.e., feedforward and feedback) between low- and higher-level sensory areas (Lamme, 2006, 2010, 2015). Representing a visual feature is necessary, but not sufficient for being conscious of it. The neural vehicle carrying that representation must also be subject to the right kind of recurrent dynamics. For instance, consciously perceiving a face consists in the feedforward activation of face selective neurons, quickly followed by a feedback signal to lower-level neurons encoding shape, color, and other visual features of the face, which in turn modulate their activity as a result (Lamme, 2010).

Especially important for our purpose is the time-window of local recurrence. Recurrent activity is typically estimated to take place between 100 and 150 milliseconds (ms) after stimulus onset (Boehler et al., 2008; Koivisto et al., 2011; Lamme & Roelfsema, 2000; Liu et al., 2009). These estimates are based on three main sources of evidence.

The retino-cortical transmission time is around 60 ms (Foxe & Simpson, 2002; Wilson et al., 1983), and signal transmission from V1 to V2 takes approximately 15 ms (Hagler et al., 2014). This implies that feedback from high- to low-level visual areas cannot happen before 75 ms after stimulus onset (Center et al., 2019). Second, visual masking, a method by which a visual stimulus is rendered subjectively invisible if it is followed by a mask after approximately 50ms, has been hypothesized by some researchers to suppress visual awareness by preventing feedback from higher- to lower-level visual areas (Boehler et al., 2008; Enns & Di Lollo, 2000; Lamme et al., 2002; Fahrenfort et al., 2007). Third, visual suppression induced by TMS to the occipital cortex around 100 ms after stimulus onset is hypothesized to prevent conscious perception through a similar mechanism (Amassian et al., 1989; Beckers & Zeki, 1995; Koivisto et al., 2017; Pascual-Leone & Walsh, 2001; Railo & Koivisto, 2012; Ro et al., 2003; Silvanto et al., 2005). Based on these results researchers hypothesize that local

² Localists would not claim, for instance, that locally recurrent processing in an isolated piece of brain tissue in a bottle would result in a conscious experience (Block, 2005, p.47). Some background conditions, such as, for example, intact thalamocortical connections, should be added for any conscious experience to occur (e.g., Alkire & Miller, 2005).

recurrence takes place between 100 and 150ms after stimulus onset, and is essentially linked to consciousness of stimuli (Lamme, 2015)³.

2. Long-lasting postdictive effects

In postdictive effects, conscious perception of a feature depends on features presented *at a later time*. For instance, in feature fusion two rapidly successive stimuli are perceived as a single entity. When a red disk is followed by a green disk after 20ms, participants report perceiving a single yellow disk, and no red or green disk at all (Figure 1a; Efron, 1967; Pilz et al., 2013). This is a postdictive effect. Both the red and green disks are required to form the yellow percept. The visual system must store the representation of the first disk until the second disk appears to integrate both representations into the percept that subjects report having. Many other postdictive effects in the range of 10-150ms have been known for decades and are well documented (Choi & Scholl, 2006; Eagleman & Sejnowski, 2000; Pilz et al., 2013; Purves et al., 1996; Schouten, 1966; Shimojo, 2014; Sugita et al., 2018)⁴.

Postdictive effects are a challenge for local theories of consciousness. Features are locally represented in the brain but the participants report that they do not see those features. In the previous example, both the green and red disks are represented in the visual system and integrated into a single yellow percept. But participants report that they don't perceive the red and green disks. Assuming that their reports can be trusted as indicating their conscious mental states (see Section 3), postdiction suggests that merely representing a feature locally does *not* lead to consciousness of that feature, which contradicts Zeki's micro-consciousness theory⁵.

³ We assume that these time estimates are approximately correct. But one should avoid drawing hasty conclusions on the links between local recurrence and consciousness from the studies just mentioned. Evidence from TMS-induced visual suppression is ambiguous, since various mechanisms could explain it without appealing to the idea that TMS specifically targets recurrent processing (Center et al., 2019; de Graaf & Sack, 2014; Mazzi et al., 2014, 2019). The same story goes for visual masking: multiple models that do not appeal to local recurrence processing can explain visual masking equally well, if not better (see e.g., Bachmann & Francis, 2013, Bridgeman, 2006; Herzog et al., 2003; Macknik, 2006; Macknik & Martinez-Conde, 2007; Pöder, 2013; Supèr & Romeo, 2012; for reviews see Bachmann & Francis, 2013; Breitmeyer & Öğmen, 2006).

⁴ One of the most documented postdictive effects is the flash-lag effect: a flashed stimulus appears to lag behind a continuously moving stimulus, even if they are actually aligned at the time of the flash (Eagleman & Sejnowski, 2000). When the moving object and the probe disappear at the same time, the flash-lag effect vanishes. When both the moving object and the probe appear at the same time, the flash-lag effect is still present. The difference between these two conditions indicates that the perception of the lag depends on what happens after the flash (i.e., whether the moving object continues to move). Therefore, the flash-lag effect is a postdictive effect (Eagleman and Sejnowski, 2000; Nijhawan, 2008; Shimojo, 2014).

⁵ Postdictive effects also occur *across* sensory modalities. This is the case in the visuo-auditory flash-lag effect (Alais & Burr, 2003; Arrighi et al., 2005; Hayashi & Murakami, 2019; Hine et al., 2003), or the motion-bounce illusion (Sekuler et al., 1997). Stiles et al. (2018) have recently demonstrated that sounds can postdictively affect the fate of visual stimuli. A sound presented *after* a visual stimulus can either cause its suppression from consciousness, or the perception of a second, illusory stimulus. Research in the domain of cross-modal postdiction is still recent. But these results could indicate, *contra* some local theories, that multisensory integration precedes consciousness.

The local recurrence theory can explain short postdiction by hypothesizing that recurrent processing kicks in at the 100ms timescale of these effects. The individual red and green disks are processed in a feedforward manner without being perceived by the subjects. Around 100ms, they are integrated into a yellow disk by recurrent processing. This is what participants consciously experience, as indicated by their reports. Although short postdictive effects constitute a problem for “extreme” local theories, they do not for more complex local views such as the local recurrence theory.

But postdictive effects also occur at much longer timescales, over 400ms. For example, Scharnowski et al. (2009) flashed two verniers with opposite offsets one after the other (a vernier stimulus is made of two vertical bars separated by a horizontal offset; Figure 1b), leading observers to report perceiving only a single fused vernier with an intermediate offset (Figure 1b). Scharnowski et al. applied TMS to the occipital cortex at different times after the onset of the first vernier. These TMS pulses strongly modulated the reported offset of the fused vernier. For early pulses, the second vernier dominated the conscious percept: observers reported perceiving an offset in the direction of the second vernier. For later TMS pulses, up to 370ms, the first vernier dominated (Figure 1b). This is a postdictive effect: TMS pulses applied 370ms *after* stimulus onset determine which percept the subjects report having.

The fact that TMS can lead to one vernier dominating the other indicates that the two verniers are represented independently for 370ms without immediately forming a single fused representation. Independent representations of each of the verniers are maintained up to 370ms. But subjects only report perceiving the postdictive outcome based on information integrated over 370ms. A subsequent study replicated similar effects (Rüter et al., 2010).

Long lasting postdictive effects also occur in the absence of TMS. Drissi-Daoudi, Doerig & Herzog (2019) used a Sequential Metacontrast paradigm (SQM; Otto et al., 2006) to study postdictive effects lasting up to 450ms. In the SQM, vertical bars are presented sequentially in different locations, eliciting the percept of a moving stream of vertical bars (Figure 1c, column 1). When a vernier is inserted in the first frame, its offset is “transported” along the stream and perceived over later frames even though these subsequent frames actually contain only straight bars (Figure 1c, column 2). When subsequent bars are straight, observers can easily report the offset direction, which indicates that the offset itself is well represented and perceived. When a vernier is inserted in a later frame, a similar effect occurs, indicating that later offsets are also well represented and perceived (Figure 1c, column 3).

In the crucial condition, *several* verniers with different offsets are inserted at different positions in the stream. These offsets integrate: observers report perceiving a *single* intermediate moving offset and cannot report on the individual offsets (Figure 1c, columns 4&5). Verniers

integrate in this way at least up to 290 ms for all observers and up to 450 ms for half of the observers. Observers report seeing only the integrated percept and cannot report on the individual verniers *even though they are separated by up to 450 ms*. In combination with the fact that the first vernier is easily seen when subsequent bars are straight, this suggests that the first vernier is represented without being perceived for up to 450 ms.

Drissi-Daoudi et al. (2019) also showed that long-lasting postdiction in the SQM is mandatory and persists when subjects attend to the stimuli. They explicitly asked subjects to report vernier offsets and to attentively focus on this simple task. In certain conditions, they also explicitly mentioned to the participants that many verniers were present, told them at which position in the SQM stream they appeared, and that they should try to perceive and report each vernier individually. The effect persisted even though the participants were aware of the details of the paradigm and were actively trying to detect the presence of the individual verniers⁶.

In a subsequent contribution, Drissi-Daoudi, Öğmen, Herzog & Cicchini (under review) studied the effects of saccades on vernier integration in the SQM. Observers fixated one degree above the center of the SQM stimulus and made a saccade to one degree below the stimulus between the presentation of the first and second vernier. Integration was mandatory in this case up to 280ms. The authors did not investigate longer durations, as they focused on trans-saccadic integration and not on the integration duration. Again, subjects represented visual features for at least 280ms without reporting being aware of those features. Long-lasting postdiction was also resistant to large changes in the retinal input caused by saccades.

Long-lasting postdictive integration in the SQM is not a small anecdotal effect. Individual vernier offset sizes are calibrated to yield around 75% discrimination performance when they are presented alone in the SQM stream. Adding a second vernier up to 450ms later brings the performance down to *chance level* (50%), a strong, qualitative change. This all-or-none behaviour suggests that SQM results can be interpreted as changes in phenomenology: either the vernier is clearly seen, or it is not seen at all. Observers do not report seeing multiple offsets. They simply report seeing a continuous stream with an offset at all times and are incapable of telling where the offsets are present in the stream, even when they are explicitly told that several verniers are presented (Otto et al., 2006).

⁶ Importantly, these effects cannot be interpreted as resulting from *criteria effects*, whereby subjects see the individual vernier offsets consciously but simply report that they do not see them because their criterion for reporting seeing the individual offsets is too conservative. Indeed, what the subjects perceived was not inferred from subjective “seen” or “not seen” reports, but from their capacity to discriminate the offsets. Subjects were simply unable to perform this discrimination task correctly. In other words, the stimuli were not only below the threshold for “subjective visibility”, as indicated for example by visibility ratings, but below the more conservative “objective visibility threshold” (Lamme, 2020), as determined by the subjects’ discrimination performance.

As another example of long lasting postdictive effects, Sun et al. (2017) showed that volition could postdictively influence the reported percept of an ambiguous visual input up to 300ms. They presented an ambiguous motion stimulus that could be perceived either as oscillating horizontally or vertically and trained participants to voluntarily choose which percept to experience. Then, they presented an auditory cue instructing which percept participants should try to perceive. Cues presented 300ms *after* the ambiguous motion stimulus onset postdictively determined which percept the subjects reported experiencing (although effect sizes are much smaller than in the previous two paradigms). This suggests that both interpretations of the ambiguous stimulus coexisted in the brain for 300ms without being consciously perceived by the subjects.

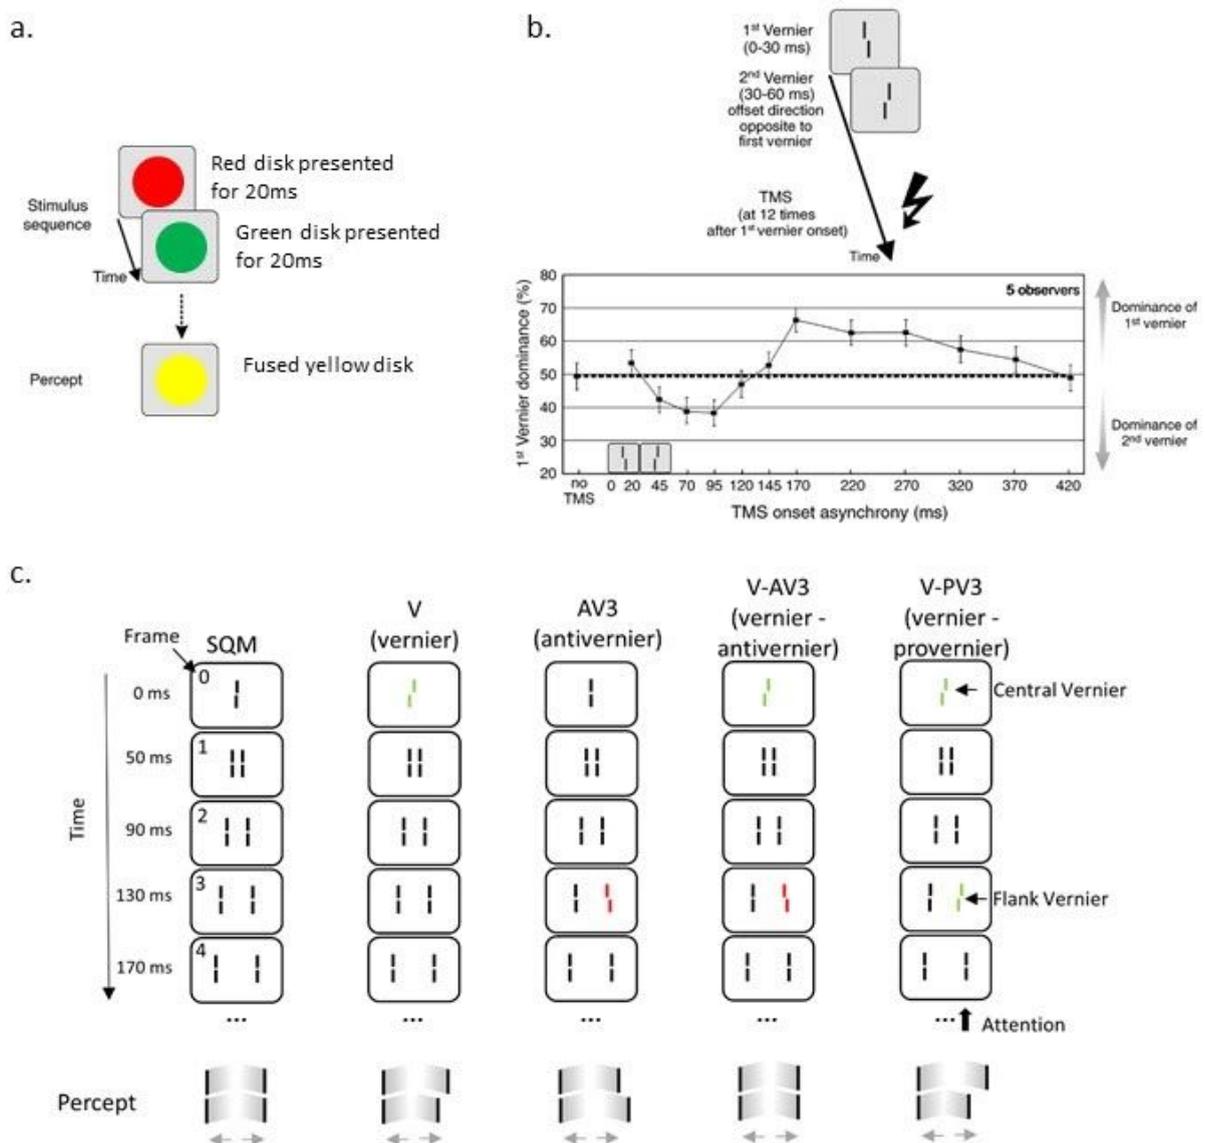


Figure 1: Postdictive effects. In postdictive effects, conscious perception of a feature depends on features presented *later*. The first feature is stably represented, but is not consciously perceived. Only the outcome of the perceptual integration between the first and second feature, presented as long as 450ms later, is perceived. **a. Color**

fusion. In this example, a red disk is presented for 20 ms, immediately followed by a 20 ms green disk at the same location. Neither the red, nor the green disk are perceived. Instead a “fused” *yellow* disk is perceived (Efron, 1967; Pilz et al., 2013). Importantly, when the red disk is presented alone for 20ms, it is easily seen. Hence, the brain represents it, but the representation is unconsciously processed until the green disk is presented and subjects experience a fused percept. **b. Long-lasting modulation of feature fusion by TMS.** Feature fusion also occurs for verniers (i.e., two vertical bars with a horizontal offset): when two verniers with opposite offsets are rapidly presented in sequence, a single fused vernier with intermediate offset is perceived. Scharnowski et al. (2009) showed that feature fusion can be modulated up to 370ms using TMS. TMS pulses were applied over the occipital cortex at different times after the onset of the first vernier to interfere with visual processing (x-axis). The y-axis shows the percentage of trials in which observers reported perceiving a fused vernier with an offset in the same direction as the first vernier. (i.e., the first vernier dominates the percept when dominance is over 50%). The TMS pulses modulated the percept up to 370ms, a long-lasting postdictive effect. This shows that both verniers are separately represented for hundreds of milliseconds without being perceived, and TMS affects the contribution to the percept of each of these representations. Figure adapted from Scharnowski et al. (2009). **c. Long-lasting postdictive effects in the SQM paradigm.** In the Sequential Metacontrast paradigm (SQM; Otto et al., 2006), sequences of vertical lines are presented, eliciting the percept of two diverging motion streams (the colors are shown here only for illustrative purposes, all the elements are the same color in experiments). When one of the vertical bars is replaced by a vernier, the vernier offset is perceived as “transported” along the motion stream, even though the vernier is presented in a single frame (columns 1-3). When several verniers are present, they are not perceived individually: the offsets integrate to yield a single percept (columns 4-5). When the offsets are in opposite directions (vernier-antivernier condition, column 4), they cancel out and no offset is perceived at all. When the offsets are in the same direction (vernier-provernier condition, column 5), they add up. This integration is mandatory: observers cannot report the individual verniers even when they know the paradigm and explicitly try to detect each vernier. Verniers as far apart as 450 ms mandatorily integrate in this way (Drissi-Daoudi et al., 2019). This indicates that the first vernier is stably represented but not consciously perceived during 450 ms – the brain stores the representation of the first vernier until the second vernier is presented and subjects experience a fused percept. Figure adapted from Drissi-Daoudi et al. (2019).

3. Why long-lasting postdiction is a problem for local theories

To the extent that one can trust the subjects’ reports as indicating consciousness, long-lasting postdiction suggests that features are stably represented in the visual system for hundreds of milliseconds, much longer than the time required for local processes such as local recurrence to happen, without the subjects being conscious of those features⁷. This is at odds with the core localist prediction.

Localists have two main options to account for long-lasting postdiction. They can deny that the local processes relevant for consciousness occur in long-lasting postdiction. Or maintain that subjects are conscious of the non-integrated features even if they report that they do not see those features. In what comes next, we analyze both alternatives in turn.

⁷ In a different debate about the discrete or continuous nature of conscious perception, Herzog et al. (2016) argue that conscious perception is constituted of *discrete* episodes resulting from the unconscious integration of sensory information. As a result, proponents of discrete views of conscious perception also hold that processing a feature does not immediately lead to conscious perception of this feature. For further discussion, see Fekete et al. (2018) and Doerig et al. (2019).

3.1. Long-lasting postdiction triggers local recurrence

One way for localists to account for long-lasting postdiction would be to deny that the adequate local processing (e.g. recurrent processing) occurs in these paradigms. Although this explanation may work in other cases, we argue that it is not cogent in long-lasting postdiction.

Sergent (2018) has recently provided an argument against localism based on studies demonstrating conscious “retro-perception” (Thibault et al., 2016; Sergent et al., 2013; Xia et al., 2016), in which an attentional cue up to 400ms after stimulus onset can trigger conscious perception of an otherwise unconsciously perceived stimulus. This could indicate that the processes happening *before* the attentional cue, including local recurrence, are not sufficient for conscious perception (Sergent, 2018).

However, because stimuli are presented at visibility threshold, whether stimuli are sufficiently strong on their own to trigger a local recurrence loop is still an open question. Instead, sensory information could be maintained unconsciously until the attentional cue triggers a reentrant loop leading to a conscious percept.

This localist explanation of the retro-perception phenomenon makes two predictions. First, there must be some mechanism capable of unconsciously maintaining information for as long as 400ms, even when that information was too weak to trigger a recurrence loop to begin with. Second, since the recurrence loop is triggered after the cue, subjects should report consciously perceiving the stimulus *after* (or at the same time as) the cue, and not *before*⁸.

Regarding the first prediction, Sergent (2018) argues that feedforward activity cannot be sufficient to maintain information for as long as 400ms. She writes:

If there is some trace to be reactivated in the posterior regions, it implies that the initial sensory processing triggered by the target itself was sufficiently strong to induce local loops in the first place, before the retro-cue. Indeed, stimuli that evoke only feed-forward activation have no lasting impact, and therefore, there would be no sensory trace to reactivate. (Sergent, 2018, p.4).

In addition, Xia et al. (2018) observed that the retro-cueing effect could be suppressed by changing the shape of the retro-cue such as to produce a masking effect through meta-contrast masking. Since visual masking is supposed to work by preventing local-recurrence, this finding indicates that local recurrence is necessary for the retro-cueing effect to occur. Which suggests that a local recurrence loop *is* initially triggered by the stimulus.

⁸ To be clear, globalism is compatible with this prediction as well. Globalists could argue that the conscious perception of the cue triggers the subsequent conscious perception of the target. But they could also hold that retro-perception is genuinely postdictive, in which case subjects unconsciously encode the stimulus, the cue, and the temporal relation between them (Trübtschek et al. 2019), and *then* come to experience the cue before the target as a single conscious episode. While globalism is compatible with both scenarios, what matters for our purpose is that localism is incompatible with the cue being perceived after the stimulus, namely, in the right temporal order.

Although current evidence indicates that it could be the case, more work is needed to support this evidence and definitely prove that recurrence is indeed triggered following the initial stimulus presentation in retro-cueing paradigms. If it is not, one alternative mechanism could be the so-called “activity-silent” short-term memory, or unconscious short-term memory (for a review, see Masse et al., 2020). Although research is still ongoing, studies have suggested that sensory information could be unconsciously maintained in short-term memory through short-term synaptic plasticity, without the need for the kind of persistent activity typically associated with maintenance through local-recurrence loops. Sensory information could be unconsciously maintained in unconscious short-term memory until the attentional cue triggers a re-activation of the sensory signals in the posterior regions, leading to a conscious percept. Although speculative, this alternative explanation cannot be ruled out at the moment.

Regarding the second prediction, there is currently no data indicating *when* subjects report perceiving the stimuli in retro-perception paradigms. If it were found that subjects report seeing the stimulus *before* the retro-cue, one would have to conclude that retro-perception is a genuinely postdictive effect. This would be inconsistent with the localist explanation: if perception of the cue triggers the recurrence loop associated with the perception of the stimulus, one should expect the stimulus to be perceived at the same time, or after the cue. As there is currently no data allowing us to decide one way or the other, we will leave retro-perception aside in the rest of this article.

Localists cannot tell the same story in the case of long-lasting postdictive effects. For instance, in the SQM paradigm, the presentation of the first vernier must be sufficient to trigger a local recurrence loop on its own, since it is easily perceived when it is not followed by a vernier with an opposite offset up to 400ms later (Drissi-Daoudi et al., 2019). As the anti-vernier is presented much later than the first vernier, it cannot retroactively prevent the first vernier from triggering a local recurrence loop. Of course, there is also no way for the perceptual system to “know” 400ms in advance that it shouldn’t trigger a recurrence loop following the presentation of the first vernier and store the representation in unconscious short-term memory instead.

As a result, local recurrence theorists have to maintain that local recurrence *does* happen from 100 to 150ms after the onset of the first vernier. Yet, participants do *not* report perceiving the first vernier when an anti-vernier is presented up to 400ms after its onset. If their reports reliably indicate their conscious percepts, local recurrence (or other local processes that require less than 400ms) is not sufficient for consciousness.

3.2. Overflow doesn't come for free

Alternatively, localists can provide the following explanation: subjects are *phenomenally* conscious of the stimuli 100 to 150ms after stimulus onset due to local recurrence. But since it takes longer to consciously access a feature than to be phenomenally conscious of it, only the integrated stimulus is available for access-consciousness (and report). By the time the percept is available for access-consciousness, the postdictive integration has already taken place, leaving the phenomenally conscious perception of the non-integrated stimuli behind. Access-consciousness overwrites phenomenal consciousness after the fact⁹.

This hypothesis is consistent with a *change blindness* interpretation of long-lasting postdiction effects: subjects do not report seeing the change between the non-integrated stimuli because the integrated percept overwrites the non-integrated stimuli, causing subjects to be “blind” to the phenomenally conscious, non-integrated stimuli.

This alternative explanation is problematic for two reasons. First, it commits local theorists to a strong and counterintuitive version of the overflow hypothesis. Second, it makes local theories particularly difficult to test, which has the downside of nullifying the weight of the evidence taken to support localism in the first place. Let us develop those two points.

The overflow hypothesis is intuitively appealing in cases like peripheral vision, when subjects do not attend to visual features, or in cases of working memory overload (e.g., Block, 1995, 2007; Bronfman et al., 2014; Fei-Fei et al., 2007; Haun et al., 2017; Sperling, 1960). These cases essentially capture what we call the *weak* overflow hypothesis: when attention is distracted (or working memory overloaded), the fact that subjects don't report seeing a visual feature, or can't report on the identity of a feature, is not a good reason for believing that they didn't have any conscious experience of that feature at all.

Contrast this with a stronger version of the overflow hypothesis: healthy subjects with nothing else to do but to detect some visual features presented right in front of them, and who systematically deny seeing those features, could nevertheless be phenomenally conscious of those features. One might have qualms with the claim that cases like these are coherent at all. But granting that such cases are possible, it is clear that the stronger version of the overflow hypothesis is much less intuitive than the weak version.

Proponents of local theories do not usually adopt this strong view. For example, Block denies that localists are committed to the counter-intuitive version of the overflow hypothesis (see Block,

⁹ See Phillips (2011) for a discussion of a similar hypothesis.

2007, p.484-485; Block, 2011b, p.446)¹⁰. And Lamme (2004) explicitly rejects it. For instance, he writes that stimuli should be considered *unconscious* when they are *attended*, and yet, *unreported*:

What exactly should we think of the distinction between conscious and unconscious (...) if it is not strictly related to our own experience of reportability? From a definition point, the distinction is simple: *unconscious stimuli or stimulus properties are those that we cannot report about, even when attended to.* (Our emphasis, Lamme, 2004, p.865)

To account for long-lasting postdiction, localists must rely on the *strong* version of the overflow hypothesis. Indeed, in long-lasting postdiction, stimuli are presented in the fovea, subjects attend to the stimuli, they are healthy, their working memory capacities are not overloaded, the reports are provided immediately after the perception of the stimuli. Still, subjects cannot report the offset direction, number or position of the non-integrated stimuli, even when they are aware of the details of the experimental paradigm and goal of the experiment (Section 2).

Moreover, the conditions in which stimuli are presented in long-lasting postdiction paradigms are different from the kind of conditions known to elicit change blindness or inattentional blindness. In these latter cases, subjects fail to detect stimuli because of attentional manipulations (Mack & Rock, 1998), or because stimuli are presented in the visual periphery (Rosenholtz, 2020). In long-lasting postdiction, stimuli are presented in the fovea, and subjects attend to them.

In addition, contrary to change blindness or inattentional blindness, long-lasting postdiction effects persist even when subjects know what they are looking for, and, in the SQM paradigm, are made aware of the position at which the relevant vernier should appear (Drissi-Daoudi et al., 2019).

For these reasons, the change blindness explanation can be maintained only if one adopts the strong overflow view: participants who attend to stimuli presented foveally can nevertheless be subject to change blindness, and thereby fail to realize that they *do* have phenomenal experiences of the non-integrated stimuli.

Localists could answer that being committed to the strong overflow hypothesis is not too worrying. True theories can have counterintuitive consequences. We agree. Merely pointing out that local theories have a counterintuitive consequence is not enough. But there is an additional problem with the strong overflow view.

Adopting the strong overflow view would make local theories particularly difficult to falsify (Cohen & Dennet, 2011; Doerig, et al., 2019; Michel, 2020). The view implies that localists can appeal to inaccessible phenomenally conscious experiences whenever it pleases them. This includes all the cases that could potentially threaten to falsify their preferred theories by suggesting that local processes can occur unconsciously (such as long-lasting postdiction cases).

¹⁰ This is not to say that Block and other localists are not at least open to the idea that the strong view *could* be true.

Localists would probably answer that if local theories do a better job at explaining the available evidence overall, we should be ready to accept the consequences of those theories, even if some of those consequences are (at the moment) untestable and counterintuitive. That is, an inference to the best explanation could lead us to accept local theories, even if some of their consequences are untestable (Block, 2007).

But the problem goes deeper. For a simple reason: however local you think consciousness is, there's always someone to believe that it is more local than you do. Extreme localists always win at the overflow game. Given the strong overflow hypothesis, it is open for proponents of more extreme local theories, such as the micro-consciousness theory, to deny unconscious perception in the cases that the champions of the local recurrence theory hold as paradigmatic cases of unconscious perception. These include visual masking, or visual suppression with TMS (see Section 1). But these phenomena are supposed to provide the main source of evidence for the local recurrence theory. As a result, proponents of the local recurrence theory would be left with no clear reason to believe that phenomenal consciousness correlates with local recurrence, which would make this hypothesis arbitrary. Adopting the strong overflow hypothesis would protect the local recurrence theory only at the price of nullifying the weight of the evidence taken to support the theory in the first place.

Localists have to provide a justification for accepting unconscious perception, rather than inaccessible phenomenal consciousness, in cases that are interpreted as supporting the theory. Here is Lamme's (2015) justification in the case of visual masking:

It is safe to assume invisibility in masking, *because there is no conceivable reason that could prevent the subject from reporting his visual percept, had he had one*: the subject is sitting there, focussing his full attention to the target location, ready to push the button as soon as he sees the target. The not-seeing can therefore not be attributed to the absence of attention, to a lapse of memory, or to any other cognitive function sitting between a potentially conscious sensation and its report. (Our emphasis, Lamme, 2015, p.34)

The reasons that Lamme (2015) invokes for holding that visual masking is a case of unconscious perception are exactly analogous to those that we put forward for maintaining that subjects are unconscious of the non-integrated stimuli in long-lasting postdiction. If the former cases count as supporting the local recurrence theory, the latter have the power to falsify it. On the other hand, if the former cases do not count as supporting the theory, there's no reason to believe that consciousness should be associated with local recurrence in the first place. Localists can't have their overflowing cake and eat it too. Either no evidence supports the theory, or long-lasting postdiction could disprove it.

Localists can save their theories only at the price of accepting the counterintuitive strong overflow view, which nullifies the weight of the evidence supposed to support those theories. On the

other hand, if localists do not want to be committed to the strong overflow view, long-lasting postdiction constitutes a new empirical challenge for local theories. Long-lasting postdictive effects show that visual features can be locally represented for hundreds of milliseconds, much longer than the time required for recurrent processing (or other local processes) to occur, without being consciously perceived. This contradicts the core prediction of local theories.

Conclusion

Theories of consciousness should account for long-lasting postdiction. In this article, we focused on the challenges that this phenomenon could constitute for local theories, because localists postulate that conscious perception happens at a relatively short timescale. But other theories, such as global workspace, or higher-order theories, could have difficulties accounting for long-lasting postdiction as well. For example, the 400ms window of perceptual integration should be reconciled with the estimated 300ms time window of global broadcast suggested by global workspace theorists (Dehaene & Changeux, 2011). Among other possibilities, this could be done by hypothesizing discrete global broadcasting episodes at the 300 to 400ms timescale (Herzog et al., 2016). Future research is needed to tell if global theories of consciousness can successfully account for long-lasting postdiction. In the meantime, we have argued that these effects constitute a serious challenge to local theories.

Acknowledgements

We thank Ned Block, Mouslim Cherkaoui, Axel Cleeremans, Michael Herzog, Victor Lamme, Hakwan Lau, Jorge Morales, Henry Railo, Claire Sergent, and Taylor Webb for helpful comments. Adrien Doerig was supported by the Swiss National Science Foundation grant ‘Basics of visual processing: from elements to figures’ (176153).

References

- Alais, D., & Burr, D. (2003). The “flash-lag” effect occurs in audition and cross-modally. *Current Biology*, *13*(1), 59–63.
- Alkire, M. T., & Miller, J. (2005). General anesthesia and the neural correlates of consciousness. *Prog. Brain Res.*, *150*, 229–244
- Amassian, V. E., Cracco, R. Q., Maccabee, P. J., Cracco, J. B., Rudell, A., & Eberle, L. (1989). Suppression of visual perception by magnetic coil stimulation of human occipital cortex. *Electroencephalography and Clinical Neurophysiology*, *74*, 458–462.
- Arrighi, R., Alais, D. & Burr, D. (2005). Neural latencies do not explain the auditory and audio-visual flash-lag effect. *Vision Research*, *45*, 2917–2925.
- Bachmann, T., & Francis, G. (2014). Visual masking: Studying perception, attention, and consciousness. *Visual Masking: Studying Perception, Attention, and Consciousness*. San Diego, CA, US: Elsevier Academic Press.
- Bagattini, C., Mazzi, C., & Savazzi, S. (2015). Waves of awareness for occipital and parietal phosphene perception. *Neuropsychologia*, *70*, 114–125.
- Bar, M. (2003). A cortical mechanism for triggering top-down facilitation in visual object recognition. *Journal of Cognitive Neuroscience*, *15*, 600–609.
- Bar, M., Kassam, K. S., Ghuman, A. S., Boshyan, J., Schmid, A. M., Dale, A. M., ... Halgren, E. (2006). Top-down facilitation of visual recognition. *Proceedings of the National Academy of Sciences of the United States of America*, *103*(2), 449 LP – 454.
- Block, N. (1995). On a Confusion about a Function of Consciousness. *Behavioral and Brain Sciences*, *18*, pp. 227–287.
- Block, N. (2007). Consciousness, Accessibility, and the Mesh Between Psychology and Neuroscience, *Behavioral and Brain Sciences*, *30*, pp. 481–499.
- Block, N. (2011a). Perceptual consciousness overflows cognitive access. *Trends in Cognitive Sciences*, *15*(12), 567–575.
- Block, N. (2011b). Response to Rosenthal and Weisberg. *Analysis*, *71*(3), 443–448.
- Block, N. (forthcoming). What is wrong with the no-report paradigm and how to fix it. *Trends in Cognitive Sciences*.
- Beckers, G., & Zeki, S. (1995). The consequences of inactivating areas V1 and V5 on visual motion perception. *Brain*, *118*(1), 49–60.
- Boehler, C.N., Schoenfeld, M.A., Heinze, H.J., Hopf, J.M., 2008. Rapid recurrent processing gates awareness in primary visual cortex. *Proc. Nat. Acad. Sci. U.S.A.* *105*, 8742–8747.
- Boly, M., Massimini, M., Tsuchiya, N., Postle, B. R., Koch, C., & Tononi, G. (2017). Are the Neural Correlates of Consciousness in the Front, or in the Back of the Cerebral Cortex? Clinical and Neuroimaging Evidence. *The Journal of Neuroscience*, *37*(40), 9603–9613.

- Bor, D., Schwartzman, D. J., Barrett, A. B., & Seth, A. K. (2017). Theta-burst transcranial magnetic stimulation to the prefrontal or parietal cortex does not impair metacognitive visual awareness. *PLoS One*, 12(2), e0171793.
- Breitmeyer, B. G., & Öğmen, H. (2006). *Visual masking: Times slices through conscious and unconscious vision*. Oxford: Oxford University Press.
- Breitmeyer, B. (2007). Visual masking: Past accomplishments, present status, future developments. *Advances in Cognitive Psychology*, 3(1–2), 9–20.
- Bridgeman, B. (1980). Temporal response characteristics of cells in monkey striate cortex measured with metacontrast masking and brightness discrimination. *Brain Research*, 196, 347–364.
- Bridgeman, B. (2006). Contributions of lateral inhibition to object substitution masking and attention. *Vision Research*, 46, 4075–4082.
- Bronfman, Z. Z., Brezis, N., Jacobson, H., & Usher, M. (2014). We See More Than We Can Report: “Cost Free” Color Phenomenality Outside Focal Attention. *Psychological Science*, 25(May), 1–10.
- Brown, R., Lau, H., & LeDoux, J. E. (2019). Understanding the Higher-Order Approach to Consciousness. *Trends in Cognitive Sciences*, 23(9), 754–768.
- Center, E. G., Knight, R., Fabiani, M., Gratton, G., & Beck, D. M. (2019). Examining the role of feedback in TMS-induced visual suppression : A cautionary tale. *Consciousness and Cognition*, 75(April), 102805.
- Choi, H., & Scholl, B. J. (2006). Perceiving causality after the fact: Postdiction in the temporal dynamics of causal perception. *Perception*, 35(3), 385–399.
- Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15(8), 358–364.
- Cohen, M. A., Dennett, D. C., & Kanwisher, N. (2016). What is the Bandwidth of Perceptual Experience ? *Trends in Cognitive Sciences*, 20(5), 324–335.
- Cohen, M. A., Cavanagh, P., Chun, M. M., & Nakayama, K. (2012). The attentional requirements of consciousness. *Trends in Cognitive Sciences*, 16(8), 411–417.
- Dehaene, S., & Changeux, J. P. (2011). Experimental and Theoretical Approaches to Conscious Processing. *Neuron*, 70(2), 200–227.
- Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10(5), 204–211.
- de Graaf, T. A., & Sack, A. T. (2014). Using brain stimulation to disentangle neural correlates of conscious vision. *Frontiers in Psychology*, 5(September), 1019.
- Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of Experimental Psychology: General*, 129(4), 481–507.
- Doerig, A., Schurger, A., Hess, K., & Herzog, M. H. (2019). The unfolding argument: Why IIT and other causal structure theories cannot explain consciousness. *Consciousness and Cognition*, 72, 49–59.

- Doerig, A., Scharnowski, F., & Herzog, M. H. (2019). Building perception block by block: a response to Fekete et al. *Neuroscience of consciousness*, 2019(1), niy012.
- Drissi-Daoudi, L., Doerig, A., & Herzog, M. H. (2019). Feature integration within discrete time windows. *Nature Communications*, 10(1), 4901.
- Drissi-Daoudi, Öğmen, H., Herzog, H., & Cicchini, M. G., (under review). Object identity determines trans-saccadic integration.
- Eagleman, D. M., & Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science*, 287(5460), 2036–2038.
- Efron, R. (1967). The duration of the present. *Annals of the New York Academy of Sciences*, 138(1), 713–729.
- Enns, J. T. & Di Lollo V. (2000). What’s new in visual masking? *Trends in Cognitive Sciences*, 4 (9), 345-352.
- Fei-Fei, L., Iyer, A., Koch, C., & Perona, P. (2007). What do we perceive in a glance of a real-world scene? *Journal of Vision*, 7(1), 1–29.
- Herzog, M. H., Ernst, U. A., Etzold, A., & Eurich, C. W. (2003). Local interactions in neural networks explain global effects in Gestalt processing and masking. *Neural Computation*, 15, 2091-2113.
- Herzog, M. H., Kammer, T., & Scharnowski, F. (2016). Time slices: what is the duration of a percept?. *PLoS biology*, 14(4), e1002433.
- Hine, T. J., White, A. M. V, & Chappell, M. (2003). Is there an auditory-visual flash-lag effect? *Clinical & Experimental Ophthalmology*, 31(3), 254–257.
- Fahrenfort, J. J., Scholte, H. S. & Lamme, V. A. F. (2007). Masking disrupts reentrant processing in human visual cortex. *Journal of Cognitive Neuroscience*, 19 (9), 1488-1497.
- Fahrenfort, J. J., Scholte, H. S., & Lamme, V. A. F. (2008). The spatiotemporal profile of cortical processing leading up to visual perception. *Journal of Vision*, 8(1):1–12.
- Fekete, T., Van de Cruys, S., Ekroll, V., & van Leeuwen, C. (2018). In the interest of saving time: a critique of discrete perception. *Neuroscience of consciousness*, 2018(1), niy003.
- Francis, G. (2000). Quantitative theories of metacontrast masking. *Psychological Review*, 107(4), 768–785.
- Foxe, J. J., & Simpson, G. V. (2002). Flow of activation from V1 to frontal cortex in humans: A framework for defining “early” visual processing. *Experimental Brain Research*, 142, 139–150.
- Hagler, D. J. (2014). Visual field asymmetries in visual evoked responses. *Journal of Vision*, 14(14), 1-19.
- Haun, A. M., Tononi, G., Koch, C., & Tsuchiya, N. (2017). Are we underestimating the richness of visual experience? *Neuroscience of Consciousness*, 3(1), 1–4.
- Hayashi, R., & Murakami, I. (2019). Distinct mechanisms of temporal binding in generalized and cross-modal flash-lag effects. *Scientific Reports*, 9(1), 1–12.
- Hochstein, S., & Ahissar, M. (2002). View from the top: Hierarchies and reverse hierarchies in the visual system. *Neuron*, 36(5), 791–804.
- Koch, C., & Tsuchiya, N. (2007). Attention and consciousness: two distinct brain processes. *Trends in Cognitive Sciences*, 11(1), 16–22.

- Koivisto, M., Railo, H., Salminen-Vaparanta, N., (2011). Transcranial magnetic stimulation of early visual cortex interferes with subjective visual awareness and objective forced-choice performance. *Conscious. Cognit.* 20, 288–298.
- Koivisto, M., Harjuniemi, I., Railo, H., Salminen-Vaparanta, N., & Revonsuo, A. (2017). Transcranial magnetic stimulation of early visual cortex suppresses conscious representations in a dichotomous manner without gradually decreasing their precision. *NeuroImage*, 158, 308–318.
- Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14(7), 301–307.
- Knotts, J. D., Odegaard, B., Lau, H., & Rosenthal, D. (2019). Subjective inflation: phenomenology’s get-rich-quick scheme. *Current Opinion in Psychology*, 29, 49–55.
- Kveraga, K., Boshyan, J., & Bar, M. (2007). Magnocellular projections as the trigger of top-down facilitation in recognition. *Journal of Neuroscience*, 27, 13232–13240.
- Lamme, V. A. F. & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences*, 23 (11), 571-9.
- Lamme, V. A. F., Zipser, K. & Spekreijse, H. (2002). Masking interrupts figure-ground signals in V1. *Journal of Cognitive Neuroscience*, 14 (7), 1044-53.
- Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1), 12–18.
- Lamme, V. A. F. (2004). Separate neural definitions of visual consciousness and visual attention; a case for phenomenal awareness. *Neural Networks*, 17(5–6), 861–872.
- Lamme, V. A. F. (2006). Towards a true neural stance on consciousness. *Trends in Cognitive Sciences*, 10(11), 494–501.
- Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, 1(3), 204–220.
- Lamme, V. A. F. (2015). The Crack of Dawn. *Open MIND*, 22:1–34.
- Landman, R., Spekreijse, H., Lamme V. A. F. (2003) Large capacity storage of integrated objects before change blindness. *Vision Research*, 43, 149–164.
- Lau, H., & Rosenthal, D. (2011). Empirical support for higher-order theories of conscious awareness. *Trends in Cognitive Sciences*, 15(8), 365–373.
- Liu, H., Agam, Y., Madsen, J.R., Kreiman, G. (2009). Timing, timing, timing: fast decoding of object information from intracranial field potentials in human visual cortex. *Neuron* 62, 281–290.
- Mack, A. and Rock, I. (1998). *Inattentional Blindness*, MIT Press.
- Mack, A., & Clarke, J. (2012). Gist perception requires attention. *Journal of Vision*, 10(7), 187–187.
- Macknik, S. L. (2006). Visual masking approaches to visual awareness. *Progress in Brain Research*, 155, 177–215.
- Macknik, S. L., & Livingstone, M. S. (1998). Neuronal correlates of visibility and invisibility in the primate

- visual system. *Nature Neuroscience*, 1, 144-149.
- Macknik, S. L., Martinez-Conde, S., & Haglund, M. M. (2000). The role of spatiotemporal edges in visibility and visual masking. *Proceedings of the National Academy of Science USA*, 97, 7556-7560.
- Macknik, S. L., & Martinez-Conde, S. (2007). The role of feedback in visual masking and visual processing. *Advances in Cognitive Psychology*, 3(1-2), 125-152.
- Macknik, S. L., & Martinez-Conde, S. (2004). Dichoptic visual masking reveals that early binocular neurons exhibit weak interocular suppression: implications for binocular vision and visual awareness. *Journal of Cognitive Neuroscience*, 16, 1-11.
- Mazzi, C., Mancini, F., & Savazzi, S. (2014). Can IPS reach visual awareness without V1? Evidence from TMS in healthy subjects and hemianopic patients. *Neuropsychologia*, 64, 134-144.
- Mazzi, C., Savazzi, S., & Silvanto, J. (2019). On the “blindness” of blindsight: What is the evidence for phenomenal awareness in the absence of primary visual cortex (V1)? *Neuropsychologia*, 128(June 2017), 103-108.
- Michel, M. (2020). Consciousness Science Underdetermined: A short history of endless debates. *Ergo*, 6(8).
- Michel, M., & Morales, J. (2019). Minority Reports: Consciousness and the Prefrontal Cortex. *Mind & Language*.
- Naccache, L. (2018). Why and how access consciousness can account for phenomenal consciousness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373.
- Nijhawan, R. (2008). Visual prediction: Psychophysics and neurophysiology of compensation for time delays. *Behavioral and Brain Sciences*, 31(2), 179-239.
- Odegaard, B., Knight, R. T., & Lau, H. (2017). Should a few null findings falsify prefrontal theories of conscious perception? *The Journal of Neuroscience*, 37(40), 9593-9602.
- Odegaard, B., Chang, M. Y., Lau, H., & Cheung, S.-H. (2018). Inflation versus filling-in: why we feel we see more than we actually do in peripheral vision. *Philosophical Transactions of the Royal Society of London. Series B, Biological Sciences*, 373(1755).
- Otto, T. U., Ögmen, H., & Herzog, M. H. (2006). The flight path of the phoenix – The visible trace of invisible elements in human vision. *Journal of Vision*, 6(10), 1079-1086.
- Papanikolaou, A., Keliris, G. A., Papageorgiou, T. D., Schiefer, U., Logothetis, N. K., & Smirnakis, S. M. (2019). Organization of area hV5/MT+ in subjects with homonymous visual field defects. *NeuroImage*, 190(March), 254-268.
- Pascual-Leone, A. & Walsh, V. (2001). Fast backprojections from the motion to the primary visual area necessary for visual awareness. *Science*, 292 (5516), 510-512.
- Phillips, I. B. (2011). Perception and iconic memory: What Sperling doesn't show. *Mind and Language*, 26(4), 381-411.
- Pilz, K. S., Zimmermann, C., Scholz, J., & Herzog, M. H. (2013). Long-lasting visual integration of form, motion, and color as revealed by visual masking. *Journal of Vision*, 13(10), 12-12.
- Pinto, Y., Otten, M., Seth, A. K., Vandembroucke, A. R. E., Sligte, I. G., & Lamme, V. A. F. (2017). Conscious Visual Memory With Minimal Attention. *Journal of Experimental Psychology: General*, 146(2),

214–226.

- Pitts, M. A., Padwal, J., Fennelly, D., Martinez, A., & Hillyard, S. A. (2014a). Gamma band activity and the P3 reflect post-perceptual processes, not visual awareness. *NeuroImage*, *101*, 337–350.
- Pitts, M. A., Metzler, S., & Hillyard, S. A. (2014b). Isolating neural correlates of conscious perception from neural correlates of reporting one's perception. *Frontiers in Psychology*, *5*(SEP), 1–16.
- Pöder, E. (2013). Attentional gating models of object substitution masking. *Journal of Experimental Psychology: General*, *142*(4), 1130–1141.
- Purves, D., Paydarfar, J. A., & Andrews, T. J. (1996). The wagon wheel illusion in movies and reality. *Proceedings of the National Academy of Sciences*, *93*(8), 3693–3697.
- Railo, H., & Koivisto, M. (2012). Two means of suppressing visual awareness: A direct comparison of visual masking and transcranial magnetic stimulation. *Cortex*, *48*(3), 333–343.
- Ro, T., Breitmeyer, B., Burton, P., Singhal, N.S., Lane, D. (2003). Feedback contributions to visual awareness in human occipital cortex. *Current Biology*, *13*(12), 1038–1041.
- Rounis, E., Maniscalco, B., Rothwell, J. C., Passingham, R. E., & Lau, H. (2010). Theta-burst transcranial magnetic stimulation to the prefrontal cortex impairs metacognitive visual awareness. *Cognitive Neuroscience*, *1*(3), 165–175.
- Rosenholtz, R. (2020). Demystifying visual awareness: Peripheral encoding plus limited decision complexity resolve the paradox of rich visual experience and curious perceptual failures. *Attention, Perception, and Psychophysics*.
- Rüter, J., Kammer, T., & Herzog, M. H. (2010). When transcranial magnetic stimulation (TMS) modulates feature integration. *European Journal of Neuroscience*, *32*(11), 1951–1958.
- Rutiku, R., Tulver, K., Aru, J., & Bachmann, T. (2016). Visual masking with frontally applied pre-stimulus TMS and its subject-specific neural correlates. *Brain Research*, *1642*, 136–145.
- Sekuler, R., Sekuler, A. B., & Lau, R. (1997). Sound alters visual motion perception. *Nature*, *385*, 308.
- Sergent, C., Wyart, V., Babo-Rebelo, M., Cohen, L., Naccache, L., & Tallon-Baudry, C. (2013). Cueing attention after the stimulus is gone can retrospectively trigger conscious perception. *Current Biology*, *23*(2), 150–155.
- Sergent, C. (2018). The offline stream of conscious representations. *Philosophical Transactions of the Royal Society B: Biological Sciences*, *373*(1755).
- Scharnowski, F., Rüter, J., Jolij, J., Hermens, F., Kammer, T., & Herzog, M. H. (2009). Long-lasting modulation of feature integration by transcranial magnetic stimulation. *Journal of Vision*, *9*(6), 1–1.
- Schouten, J. F. (1966). Subjective stroboscopy and a model of visual movement detectors. *IPO Annual Progress Report*, *1*, 78–80.
- Shimojo, S. (2014). Postdiction: Its implications on visual awareness, hindsight, and sense of agency. *Frontiers in Psychology*, *5*(MAR), 1–19.
- Silvanto, J., Lavie, N. & Walsh, V. (2005). Double dissociation of V1 and V5/MT activity in visual awareness. *Cerebral Cortex*, *15* (11), 1736-1741.

- Silvanto, J. (2014). Is primary visual cortex necessary for visual awareness? *Trends in Neurosciences*, 37(11), 618–619.
- Sligte, I. G., Vandenbroucke, A. R. E., Scholte, H. S., & Lamme, V. A. F. (2010). Detailed sensory memory, sloppy working memory. *Frontiers in Psychology*, 1, 12.
- Sperling, G. (1960). The Information Available in Brief Visual Presentations. *Psychological Monographs: General and Applied*, 74(11).
- Stiles, N. R. B., Li, M., Levitan, C. A., Kamitani, Y., & Shimojo, S. (2018). What you saw is what you will hear: Two new illusions with audiovisual postdictive effects. *PLoS ONE*, 13(10), 1–22.
- Sugita, Y., Hidaka, S., & Teramoto, W. (2018). Visual percepts modify iconic memory in humans. *Scientific Reports*, 8(1), 1–2.
- Sun, L., Frank, S. M., Hartstein, K. C., Hassan, W., & Tse, P. U. (2017). Back from the future: Volitional postdiction of perceived apparent motion direction. *Vision Research*, 140, 133–139.
- Supèr, H., & Romeo, A. (2012). Masking of figure-ground texture and single targets by surround inhibition: A computational spiking model. *PLoS ONE*, 7(2), 1–10.
- Tapia, E., & Breitmeyer, B. G. (2011). Visual consciousness revisited. *Psychological Science*, 22(7), 934–942.
- Thibault, L., Van Den Berg, R., Cavanagh, P., & Sergent, C. (2016). Retrospective attention gates discrete conscious access to past sensory stimuli. *PLoS ONE*, 11(2), 1–13.
- Trübutschek, D., Marti, S., Ueberschär, H., & Dehaene, S. (2019). Probing the limits of activity-silent non-conscious working memory. *Proceedings of the National Academy of Sciences of the United States of America*, 116(28), 14358–14367.
- Tsuchiya, N., Wilke, M., Frässle, S., & Lamme, V. A. F. (2015). No-Report Paradigms: Extracting the True Neural Correlates of Consciousness. *Trends in Cognitive Sciences*, 19(12), 757–770.
- Thielscher, A., Reichenbach, A., Uurbil, K., & Uluda, K. (2010). The cortical site of visual suppression by transcranial magnetic stimulation. *Cerebral Cortex*, 20(2), 328–338.
- Xia, Y., Morimoto, Y., & Noguchi, Y. (2016). Retrospective triggering of conscious perception by an interstimulus interaction. *Journal of Vision*, 16(7), 1–8.
- Ward, E. J., Bear, A., & Scholl, B. J. (2016). Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining whether awareness overflows access. *Cognition*, 152, 78–86.
- Ward, E. J. (2018). Downgraded phenomenology : how conscious overflow lost its richness. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 373.
- Wilson, C. L., Babb, T. L., Halgren, E., & Crandall, P. H. (1983). Visual Receptive Fields And Response Properties Of Neurons In Human Temporal Lobe And Visual Pathways. *Brain*, 106(2), 473–502.
- Zeki, S., & Bartels, A. (1999). Toward a Theory of Visual Consciousness. *Consciousness and Cognition*, 8(2), 225–259.
- Zeki, S. (2003). The disunity of consciousness. *Trends in Cognitive Sciences*, 7(5), 214–218.