Check for updates

# Does Perceptual Consciousness Overflow Cognitive Access? The Challenge from Probabilistic, Hierarchical Processes

STEVEN GROSS AND JONATHAN FLOMBAUM

**Abstract:** Does perceptual consciousness require cognitive access? Ned Block argues that it does not. Central to his case are visual memory experiments that employ post-stimulus cueing—in particular, Sperling's classic partial report studies, change-detection work by Lamme and colleagues, and a recent paper by Bronfman and colleagues that exploits our perception of 'gist' properties. We argue *contra* Block that these experiments do not support his claim. Our reinterpretations differ from previous critics' in challenging as well a longstanding and common view of visual memory as involving declining capacity across a series of stores. We conclude by discussing the relation of probabilistic perceptual representations and phenomenal consciousness.

## 1. Introduction

Does perceptual consciousness require cognitive access? Correlatively, do the core neural correlates of perceptual consciousness include neural regions or processes implicated specifically in cognitive access? In an important series of papers, Ned Block (1995, 2007a, 2007b, 2008, 2011, 2014a) argues that they do not: for example, one can have a visual experience as of some letter even though this content does not reach the mechanisms that support reasoning, planning, and the like. In that sense, one can have a conscious perception without being able to report or think about it. Central to Block's case are visual memory experiments that employ post-stimulus cueing—in particular, Sperling's (1960) classic partial report studies, change-detection work by Lamme and colleagues (e.g. Vandenbroucke *et al.*, 2011), and a recent paper by Bronfman *et al.* (2014) that exploits our perception of 'gist' properties. We argue *contra* Block (also Burge, 2007; Dretske, 2006; and Tye, 2006) that these experiments do not support his claim. Our response differs from previous critics' in challenging as well a longstanding and common view of visual memory as involving a series of stores of declining capacity. This view, we suggest, does not

**Acknowledgements:** Thanks to Justin Bledin, Peter Carruthers, Krish Eswaran, Alison Gopnik, Hilla Jacobson, John Morrison, Ian Phillips, Brian Scholl, Barry Smith, Paul Smolensky, Marius Usher, an anonymous referee, and audiences at SPP 2014, SEP 2015, Johns Hopkins, University of London, and the Institute of Philosophy of Mind and Cognition (National Yang Ming University, Taiwan). Special thanks to Ned Block.

**Address for correspondence:** Department of Philosophy, Johns Hopkins University, Baltimore, MD 21218, USA.
**Email:** sgross11@jhu.edu

*Mind & Language*, Vol. 32, No. 3 June 2017, pp. 358–391.
*© 2017 John Wiley & Sons Ltd*

square with recent empirical work that emphasizes the probabilistic nature of the computations and representations perceivers deploy in response to noisy signals.

Post-stimulus cueing involves the presentation of a cue after a stimulus has gone away. The presence, timing and features of such a cue have been shown to affect participants' performance on tasks dependent on their perception of the preceding stimulus—even though the cue appears *after* stimulus offset. For example, Sperling (1960) presented participants with a 3x4 array of letters (also numbers in some experiments) for a variable exposure of 15 to 500 ms. In the 'full report' condition, participants were asked to report all of the letters and their locations, guessing when not certain. They accurately reported about four to five letters, on average. In the 'partial report' condition, after the array disappeared, a tone indicated to participants from which of the three rows they should report letters. When the tone occurred 300 ms after stimulus offset, they reported on average three to four letters accurately. This is a rather impressive performance in light of the 'full report' results: how could participants, who could accurately report only four to five of 12 letters when there was *no* cue, accurately report (almost) the entirety of *any* row without knowing in advance which row would be cued after offset?

Sperling argued that this is evidence of an early, sensory memory store (dubbed 'iconic memory' in Neisser, 1967) with greater capacity than visual short-term memory (also known as visual working memory). Iconic memory may have a letter capacity as large as nine to 12 letters (or larger), from which only three to five letters can be passed into short-term memory. This would account for participants' limited performance in the 'full report' condition. As for the cue, it facilitates the transfer of only the relevant row of letters into the later-stage visual memory that's crucial for report. Thus can the cue affect performance despite occurring after stimulus offset: the cue does not affect the encoding of content into the earlier store, but only what's made available for various tasks. The view that visual memory involves successive memory stores of declining capacity managed by an attentional gate-keeper—with visual working memory having a capacity of roughly four letters—is now a long-standing and common view in the field (Coltheart, 1980; Cowan, 2001).

Block (1995, 2007a, 2007b, 2008, 2011) argues that we should associate visual iconic memory—the early, high capacity stage—with visual *consciousness*. Among other things, this accommodates the report of some participants that they seemed to see all the letters, even though they couldn't report them all. If one grants this association, and if presence in later, short-term memory is necessary for cognitive access, we have Block's thesis that perceptual consciousness 'overflows'—has a greater capacity than—cognitive access. Roughly, we see more than we can say. In this case, we see all, or almost all, the letters, but can report only three to four of them.

We can bring Block's views into sharper focus and underscore their significance by comparing them to an opposed position motivated by other experimental results. In a variety of conditions, observers purport to have detailed experience of a visual scene, yet fail to notice large changes, including the appearances and disappearances of whole objects (Simons and Levin, 1997; Simons and Rensink, 2005). These and

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

other well-known 'blindness' phenomena suggest to some that our visual experience is less rich in its specific representational content than it might be natural to suppose.[1] On this view, we are subject to a kind of illusion: because we at least implicitly know that we can readily perceptually retrieve information wherever we direct our attention, we believe that we already perceive it. As Dehaene and colleagues (2006) develop the view, consciousness not only requires attention, but cognitive access: we are perceptually conscious only of what gets globally broadcast for reasoning and guidance of behaviour more generally. The core neural correlates of consciousness thus include neural areas and processes specifically implicated in cognitive access.

Block's claims are precisely the opposite. Perceptual consciousness does not require cognitive access and thus the core neural correlates of perceptual consciousness do not include neural regions or processes implicated specifically in cognitive access. Perceptual consciousness is not sparser than one might have naturally thought, but rather richer in specific content than the mechanisms of working memory that yield cognitive access. Regarding blindness phenomena and the like, we should not confuse a failure to access what we see—required for detecting and reporting a specific change—with a failure to consciously perceive it. Rather than posit an illusion, we should posit limited access.[2]

We shall argue that post-stimulus cueing results in fact do not favour Block's views. Block's interpretation of these results has been challenged before. However, previous critics have focused on Block's claim that the contents of the earlier, higher-capacity store are conscious (Phillips, 2011a, 2011b; Kouider *et al.*, 2010; Stazicker, 2011; Cohen and Dennett, 2011). They propose instead that the cue, in selectively transferring contents to the later store, can have the effect of *raising* contents to consciousness. They thus retain the idea of successive memory stores of declining capacity, disagreeing with Block rather over *which* store is associated with perceptual consciousness.[3]

Our challenge differs. We argue that post-stimulus cueing results can be explained, not only without adverting to Block's 'overflow' hypothesis, but without an assumption of declining capacity among successive memory stores and, more

[1] For striking recent work, see Freeman and Simoncelli (2011), discussed in Block (2013). They created visual metamers from photographs of real world scenes. Fixating experimental observers could not reliably determine whether an image was the crisp original or a derivation severely distorted outside the fixation point.

[2] The illusion at issue would be a *cognitive* illusion, concerning what we're disposed to *judge* (in this case, about perception). Block's claim is not that his view should be preferred simply because it does not posit an illusion of this sort. If this were Block's claim, one might counter that his own view is so committed as well, at least insofar as we're disposed to think that we are aware of all our conscious states. (Of course, it's common for scientific theories to conflict in some manner with what we find natural and intuitive.) Block (2007a) does suggest, however, some reasons to question the *specific* illusion that Dehaene *et al.* posit.

[3] Block (2011, p. 568) represents Phillips (2011a) as holding that the cue has the effect of raising unconscious representations to consciousness *within* iconic memory. We don't believe this is Phillips' position, but nothing in our critique hangs on this.

© *2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

specifically, without any particular assumption concerning visual working memory's capacity limit at all. Our challenge thus extends beyond Block's overflow hypothesis to the common view of visual memory embraced equally by Block and his previous critics—not to mention many vision and memory scientists.

Our reinterpretations draw upon a growing body of work that emphasizes the probabilistic nature of the computations and representations involved in a perceiver's attempts to 'infer' the distal scene from noisy signals and then store the representations it constructs. (See Duncan, 1980, and Navon, 1984, for precursors, with recent empirical support from Bae and Flombaum, 2013; Orhan and Jacobs, 2013; and Tsal and Benoni, 2010, and theoretical support from Orhan and Ma, 2015.) On the picture we shall favour here—though, as will be clear, not all of our points require all aspects of it—perceivers construct from noisy transduced signals probabilistic representations (assignments of credences over a space of possibilities concerning the distal scene) that take into account, as best they can, expected relationships among the scene's various features; performance, in response to a specific query, then involves 'sampling' from the probabilistic representations stored in visual memory (Vul *et al.*, 2009). For example, presented with an array of letters, the visual system attempts to recover what information it can from the resulting noisy transduced signal, taking into account expected relations among lower-level features (oriented edges, curvatures, crossings, etc.) and higher-level features (letters); asked what letters are there, one returns a specific reply (when one can), with a probability that reflects the stored distribution over possibilities. From this perspective, we argue that the three-to-five item performance limit common to many visual memory tasks reflects the challenge of recovering confident representations from noisy signals in the first place; performance on post-stimulus cueing tasks is then better explained by the cues' effects on these computations—and thus on which stimuli get confidently represented—than by declining capacity of successive memory stores.

Arguments that post-stimulus cueing results do not support Block are not *per se* arguments in favor of Dehaene's view of perceptual consciousness. Indeed, much of our discussion proceeds without mentioning perceptual consciousness at all: our point is that participants' performance can be explained without reference to consciousness. After developing our challenge by considering the Sperling, Lamme, and Bronfman post-stimulus cueing results in turn, we briefly return in our conclusion to some questions that our alternative conception of visual memory poses for understanding visual consciousness. In the course of doing so, we raise the possibility that visual memory not only may not involve successive stores of declining capacity: it may not involve successive stores at all. First, however, we provide some preliminary clarifications.

## 2. Preliminaries

Before turning to our reinterpretations of the post-stimulus cueing results, we provide a few clarifications.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

First, above we spoke of capacity in terms of the relative richness or sparseness of specific representational content. In these debates, content has been deemed specific in two ways: as opposed to generic (Block, 2011, p. 568) and as opposed to general (Block, 2007b, p. 531). The former concerns a distinction between the representation of a determinable and the representation of a determinate. A determinable is roughly a property something can have only if it has another property that is a more specific way of having the first property. For example, something is coloured only if it is also some specific colour—say, red. Thus, a representation as of the letter 'A' is more specific than a representation as of some letter or other (without any specific letter specified). Although something can't *be* a letter without being a specific letter, one might allow that one can *represent* something as being a letter without representing it as being any specific letter—indeed, that one can have *perceptual* representation of this sort. Block (1983, 2007b, p. 533) agrees that one should reject a general Humean principle concerning the perceptual representation of determinables and determinates that suggests otherwise, though it seems he's willing to entertain such a requirement restricted to certain cases: 'generic conscious representations of non-square rectangles that do not specify between horizontal and vertical orientations [are] difficult to accept' (Block, 2011, p. 574). (For Hume, cf. 1740/1978, I.1.vii and Stazicker, 2011.) If one can perceptually represent something generically while not specifically, then one can be 'blind' to some aspect of a scene at some level of specificity without it being the case that one sees nothing there at all. This is crucial for rendering plausible claims of blindness and illusion outside of focal attention. The second distinction—specific (perhaps better: singular) versus general—concerns attributing a property *to some object* versus merely saying that the property is instantiated: *that* is F versus *there is* an F.[4] Capacity comparisons concern the number of specific representations in this second sense, but where arguably the most natural comparison is among representations at the same level of specificity in the first sense (though cf. Block, 2007b in reply to Kouider *et al.,* 2007).

Second, it might be objected that representations as of specific letters are conceptual and thus already beyond perception. But talk of letters here could be replaced with talk of (letter-like) shapes at a higher level of organization than some lower-level features. This reply assumes that the objection is not to perceptual representation-*as* in general, but just to perceptual representation as of *specific letters*. Block (2014b, 2015), following Burge (2010), allows non-conceptual representation-*as* in perception, distinguishing perceptual attribution and conceptual predication. It is then an empirical question whether there are perceptual attributives for specific letters. To object that representations as of specific letters are conceptual, one must either reject generally Burge's idea of perceptual attributives or argue, on empirical grounds, that representations as of specific letters are not among them.

[4] Perhaps one might try to collapse the two distinctions by analysing the latter as a certain *kind* of determinable/determinate relation.

© *2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Third, the kind of consciousness at issue is *phenomenal* consciousness: there is 'something it is like' to be in a conscious perceptual state. As is well known, Block (1995, 2005) distinguishes phenomenal consciousness and access consciousness (which, despite the name, requires access*ibility* as opposed to actual cognitive access), and he has argued that these notions are sometimes conflated. But this is not the case among Block and his critics here, who rightly take themselves to be disagreeing over a substantive matter. Where Dehaene, for example, departs from Block is over the dissociability of phenomenality from access, even granting the conceptual distinction.[5]

Finally, Block (2007a, 2007b, 2008, 2011) also uses post-stimulus cueing results to defend a methodological point. Sceptics concerning consciousness *sans* access, or at least its amenability to empirical investigation, argue that *theorists* have no access to consciousness except via *subjects*' cognitive access (e.g. for report) (Dehaene *et al.* 2006; Cohen and Dennett, 2011; Kouider *et al.,* 2010, 2012). Block replies that his interpretation of results such as Sperling's shows how inference to the best explanation can give theorists indirect access to conscious states, even absent subjects' cognitive access. With this methodological point, we fully agree—even as we demur from Block's particular interpretation of Sperling's results.

## 3. Sperling Reinterpreted

Phillips identifies a crucial assumption on which Block's interpretation of Sperling's results relies: 'Any aspect of [perceptual consciousness] present in a partial report condition would have been present even if some other partial report had been cued' (Phillips, 2011a, p. 385). If we grant that participants' reports are presumptively indicative of the contents of perceptual consciousness, then the crucial assumption permits one to determine how many letters they consciously perceive: the average number of letters participants accurately report in a cued row can simply be multiplied by the number of rows in the stimuli (cf. Sperling, 1960, p. 7—albeit without reference to consciousness).

Previous responses challenge this assumption by rejecting the claim that the cue enables a subset of *conscious* representations to pass into short-term memory, proposing instead that it selectively raises *unconscious* representations to the level of

[5] In Block (1995), Sperling's results are used to argue for the dissociability of p(henomenal)-consciousness and a(ccess)-consciousness. But, as noted in Block (2011), if being poised to enter working memory suffices for being cognitively access*ible*, then representations in visual iconic memory are cognitively accessible: they are just not all in fact cognitively *accessed* and, owing to working memory's limited capacity, *cannot* be accessed *all at once*. (Block, 2011, thus simply deploys the p-consciousness/a-consciousness distinction without adverting to Sperling's results in its defence.) One may reasonably define a-consciousness otherwise (cf. Block, 2008), but following Block (2007, p. 489) we won't worry about the word.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

consciousness (Phillips, 2011a, 2011b; Kouider *et al.*, 2010; Stazicker, 2011). They then need to explain how a cue could have this effect *after* stimulus offset and to explain away participants' reports of seeming to see all the letters. Our alternative bears these burdens as well, and we will draw on—but also improve upon—this part of their views.

Regarding the first, Phillips (2011a, 2011b) adverts to postdiction as evidence that conscious perceptual representations are constructed on the basis of information integrated over some temporal span. In postdiction, presentation of a second stimulus within 250 to 300 ms of a first affects perception of the first (Choi and Scholl, 2006; cf. also Sergent *et al.*, 2011 and 2013). For example, Sekuler *et al.* (1997) demonstrate the occurrence of sound-induced postdictive visual bounce. Two solid circles traveling towards one another along a straight line will typically appear, after convergence, to pass through one another. But an appropriate tone causes them to appear instead to bounce off each other. (The visual stimulus is ambiguous between the two.) Crucially, one gets this effect even when the tone is produced as late as 150 ms after contact would have taken place.[6] Though not mentioned by Phillips, the related literature on prior entry—the misperception of attended stimuli as having onset prior to unattended stimuli that actually onset first (e.g. Shore *et al.*, 2001 and Titchener, 1908)—likewise supports his suggestion that perceptual moments integrate over time.

As for participants' reports of seeming to see all the letters, critics have attempted to explain them away by questioning how much weight should be given to introspective reports generally, or at least in such conditions, and by noting participants' (and theorists') possible failure to distinguish specific representations as of each letter and a non-specific representation as of (generic) letters (e.g. Byrne *et al.*, 2007; Grush, 2007; Papineau, 2007—Block, 2007b replies). It's notable also that participants require a large number of practice trials for their performance to reach statistical significance (Chow, 1985), which plausibly affects participants' *expectations* concerning the stimuli. Thus, Kouider *et al.* (2010), adverting to results in de Gardelle *et al.* (2009), argue that participants' prior expectations of letters lead them both to fill in fragmentary feature representations and to generate 'letter tags'—representations as of a letter, albeit not as of any specific letter—in Sperling-like arrays (though see Block, 2011 for a reply). These various considerations are consonant with Dehaene's suggestion that participants are subject to illusions concerning what they see outside focal attention. Indeed, they may be subject both to perceptual illusions (e.g. seeing specific letters that are not there) and cognitive illusions (e.g. judging, on the basis of one's non-specific perceptual representations as of generic letters, that one sees more specific letters than one does).[7]

[6] Note that sound-induced visual bounce is cross-modal in the same way as Sperling's original experiments, which cued rows in the partial-report condition by tones (high, medium or low) that subjects had been trained to associate with the visually presented rows.

[7] Kouider *et al.* (2010) present an account of perception as probabilistic and hierarchical in many ways congenial to ours, but only to bolster their view that participants in Sperling experiments

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Although Block's critics reject the assumption that Phillips identifies, there is, as we've noted, an even more basic assumption that both Block and his critics accept—viz., that the cue selectively transfers representations from higher-capacity iconic memory to lower-capacity working memory, without affecting what gets represented in iconic memory in the first place. Block indeed maintains that one *must* accept this assumption (cf. Phillips 2011a, p. 401):

> A number of the critiques … challenge the premise that there are more than 4 items of specific phenomenology before the cue. It is important to recognize that the objectors have to agree that before the cue, there are *specific* (not just generic) visual representations of all or almost all of the … items in the Sperling … experiments. There have to be such specific representations given that any location can be cued with high accuracy of response. The locus of controversy is *whether those specific representations are phenomenal*. (Block, 2007, p. 531, original italics—the elided material mid-quote extends the claim to experiments from Lamme's group.)

Block's critics thus preserve the idea of an early, higher-capacity store, with the cue modulating which representations reach the later store. They reject only Block's claim that the contents of the earlier store are conscious.[8]

We challenge this more basic assumption as well. On our view, participants' performance limitations on visual memory tasks—including the limited number of letters reported in Sperling's 'full report' condition—reflect the difficulty of generating accurate representations from noisy, ambiguous signals, not the (declining) capacity limits of memory stores. Moreover, the computations involved in generating these representations provide ample opportunity for the cue to affect what gets

---

are subject to illusion, not to explain the effects of the cue or to question the need for rich, specific unconscious representations, as we do below.

[8] One might question the attribution to Kouider *et al*. (2010). While their Figure 1 seems clearly to depict specific unconscious letter representations in visual iconic memory, no clear commitment is found in their text. Perhaps then their figure is misleading, and its letters are meant rather to depict the distal stimulus? But Kouider and Dehaene (2007, p. 870) write:

> … preconscious processes potentially carry enough activation for conscious access, but are temporarily buffered in a non-conscious store owing to a lack of top-down attentional amplification (for instance, owing to transient occupancy of the central workspace system) … they are clearly maintained in a sensory buffer for a few hundreds of milliseconds, since they may ultimately achieve conscious access once the central workspace is freed.

Cf. Dehaene *et al*. (2006). In any event, Block (2011) certainly reads Kouider *et al*. (2010) as hypothesizing these unconscious letter representations. For instance, he quotes to this effect Cohen and Dennett (2011) who advert to Kouider *et al*. (2010) in claiming that '[p]articipants can identify cued items because their identities are stored unconsciously until the cue brings them to the focus of attention' (Cohen and Dennett, 2011, p. 359, quoted at Block, 2011, p. 568—see also Block's Figure 2a). Moreover, Kouider *et al*.'s (2012) reply to Block is most naturally read as confirming this reading.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

represented (consciously or unconsciously) in the first place; so, the cue's role need not involve the selective transfer of representations one already has. In reinterpreting Sperling's results, we thus focus on the claim that participants represent almost all the letters (and thus more than the commonly assumed capacity of visual working memory). In transitioning to Lamme's results, we target the common assumption concerning visual working memory capacity itself.

Because our proposal, developed below, does not posit highly detailed unconscious representations in iconic memory, it is immune to one of Block's main replies to previous critics—viz. that such representations lack independent motivation or evidence (Block, 2011). On the contrary, we draw upon well-supported, mainstream perception research. Likewise, we can appeal to postdiction while avoiding an uncomfortable question faced by Phillips. A plausible rationale for the temporally extended integration underlying postdiction is that it allows 'the visual system [to] take into account information from the immediate future before committing to an interpretation of the event' (Rao *et al.*, 2001, p. 1245—quoted by Phillips, 2011a, p. 391, fn. 11). But why, on Phillips' view, should this rationale apply only to conscious perceptions and not the rich, specific unconscious representations he posits? If it's a good strategy, then the visual system should similarly take the immediate future into account in the generation of the representations themselves, and not just in deciding which to render conscious—but then we should not assume that there is sufficient time before the cue for rich, specific unconscious representations to be formed.

Our alternative proceeds from the well-known fact that transduced signals—for example, patterns of firing immediately off the retina—both greatly underdetermine their distal causes and are noisy in that the same causes don't always yield the same signal. Recovering information regarding the distal scene is thus a complex and daunting task. In the case of vision, an array of light intensities must be 'unmixed' into spatially-arranged objects antecedently unknown both in number and properties (Orhan and Ma, 2015; Ma and Flombaum, 2013). This feat is accomplished by transitions, from transduced signals to representations of the distal scene, that in some manner take into account prior expectations about the scene (e.g. Stocker and Simoncelli, 2006; Girshick *et al.,* 2011), as well as expectations about relationships among features in scenes generally (e.g. Dawson, 1991; Ullman, 1979, and Orhan and Jacobs, 2013).[9] For example, representations

[9] We speak of transitions here in deference to (1) those who restrict talk of *computations* to transitions among representations and deny that transduced signals are representations and (2) those who restrict talk of *inference* to certain transitions among *conceptual* representations. But we are sometimes lax in what follows. (Burge, 2010, restricts talk of inference in this way and also denies that the retinal signal is a representation, even if it co-varies with or otherwise carries information about the distal environment. On his view, the mark of perceptual representation is rather being the output of a mechanism that sustains a perceptual constancy. Block (2014b, 2015) expresses general sympathy with Burge's outlook. From this point of view, one would reject such claims as that the retina has a *representational* capacity greater than four—cf. Block, 2007a, p. 439; though note his scare-quotes.)

© *2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

concerning 'higher-level' features can depend on representations concerning 'lower-level' features (cf. Block, 2014b)—thus nearby colinear edges (lower-level) are deemed more likely to belong to the same larger contour (higher-level) than not (Geisler *et al.,* 2001). In more sophisticated models, the dependency runs in both directions, so that lower-level representations can depend on (expectations concerning) higher-level representations as well (cf. Yuille and Kersten, 2006, on perception as 'analysis by synthesis')—and the representations at all levels can be probabilistic. Further, representations at the same level are typically not independent of one another, but rather are constrained by expectations concerning relations among them—for example, whether a contour is assigned to an object or shadow depends on other contour assignments (Geisler *et al.*, 2001).

The transitions are not guaranteed to generate accurate representations; indeed, we should expect the signal's noise and ambiguity to lead to some inaccurate results. We should expect as well that perceptual processing will fail to generate representational content at all regarding some aspects of the distal scene. Obviously, there are limits to perceptual range and acuity. But, more specifically, the transitions might function to favour representations regarding some aspects of the distal scene over others in a task-specific way; and, even taking that into account, the complexity of the task might set a bound to what they can accomplish. Consider, in particular, an increase in complexity owing to an increase in number of objects in the stimulus. More objects introduce more sources of noise as well as more features whose dependencies on one another must be accommodated, with a geometrically expanding space of possible conclusions. We can expect that there will be a limit to the number of objects and features a perceiver can accurately recover. At the limit, this would be reflected in participants' performance on relevant tasks. What is this limit? On our proposal, participants in typical visual memory experiments can report only three to five objects correctly because that's all their visual system can *confidently* recover, on average.

Let's illustrate this possibility by considering Sperling's task. What holds of perception generally holds of letter perception in particular, as is reflected in the noisy, context-dependent nature of various well-known models of letter perception (see, e.g., Grainger *et al.*, 2008, on Pandemonium; McClelland and Rumelhart, 1981, and Rumelhart and McClelland, 1982, on neural network implementations; and Norris and Kinoshita, 2012, for a recent probabilistic approach). Letter perception is hierarchical in that representations pertaining to being such-and-such a letter depend on representations concerning such lower-level features as line segments, orientations and crossings (Grainger *et al.*, 2008).[10] The challenge of arriving at letters

[10] Although hierarchical representational dependencies are relevant to the Sperling case, our more general argument does not turn on this feature. If it did, it would be open to the objection that partial report superiority might be found with ground-level representations (representations that are not dependent on any lower-level representations but are rather generated from non-representational inputs). Indeed, it's plausible that such superiority would be found, and arguably results involving colour stimuli provide examples (e.g. Makovski *et al.*, 2008).

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

from features becomes geometrically more difficult as the inclusion of additional letters adds more features to the stimulus. Idealising, imagine—for example—that the visual system is presented with two letters and extracts information concerning the probabilities of various line segments being present, their orientations, their spatial locations and relations to one another, etc. The information is noisy, so it's not *obvious* whether it's more likely that what caused the transduced signal from which this information has been extracted is 'MN' or 'NM' or any of various other possible letter pairs. Nonetheless, perhaps the visual system has enough to go on to generate a reasonable response. But now suppose we multiply by a factor of six the number of letters presented. If we limit stimuli to the 20 letters Sperling used, we have gone from $20^2$ to $20^{12}$ possibilities, albeit constrained by whatever information the visual system may have regarding, e.g., likely relations among features. Moreover, the number of lower-level features that must be extracted and retained from the noisy signal—to be used in arriving at a hypothesis regarding the distal stimulus—has also significantly increased.

How might vision deal with this? *One* possibility consistent with recent work (Orhan and Ma, 2015) is that, without selecting only a subset of features to analyse, the glut of possibilities would prevent the visual system from settling upon sensible conclusions. (Perhaps the probabilities associated with various hypotheses concerning the letters present would be too low to cross a reasonable threshold.) If so, then if one can identify letters at all, it must be because one's visual system restricts itself to something more manageable. On Sperling's 'full report' trials, a sub-region might be selected that affords good results for about three to five letters—presumably a sub-region centered on the fixation point, as is indeed suggested by the letters participants tend to produce.

Block and his critics should agree that the challenges and noisy nature of perceptual processing can limit performance. But, because they instead explain performance in Sperling's experiment by appeal to a capacity-limited store, they are committed to maintaining that the recovery limitations just emphasized have minimal impact on such tasks, only manifesting themselves more fully when perceptual systems face even greater challenges. We see no reason to make this assumption and plenty of reason to doubt it. For example, according to Bouma's Law, gathering equally good information further from central fixation requires that the periphery contain larger letters, more widely spaced (Bouma, 1970). But the stimuli were not designed accordingly in Sperling's experiments—and, even if participants had failed to fixate, the stimulus duration would have at most allowed for one saccade. Further grounds in our favour will be presented below.

What of the effect of the cue in Sperling's 'partial report' trials? A cue is itself an element of the distal scene and thus affects the transduced signal and therefore the transitions that generate representations concerning that scene. The dependencies remarked upon above offer various opportunities for a cue to affect what gets represented, including by in effect restricting or weighting inputs to perceptual computations at various stages—especially when the cue is expected to carry task-specific information. Consider again the possibility that, in order to recover any letters at all,

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

the visual system might select a subset of features to analyse. In the 'partial report' condition, a cue in effect tells the observer which locations are task-relevant (participants after all receive significant training to just this end). The observer, or her visual system, may thus do just what one does on 'full report' trials: restrict her computations to a manageable sub-region, but now to the sub-region cued as task-relevant.

There are indeed various, not necessarily mutually exclusive ways, consistent with current work on vision and memory, that one could develop our alternative proposal. We provide a few examples. In each case, as above, we advert to mainstream ideas in perception research and indicate how in principle one might plausibly apply them to the Sperling task. It's a further matter to empirically *confirm* that they play a significant role specifically in such tasks—but the same holds for *dis*confirmation.

Suppose the dependencies of higher-level features on lower-level features are associated with a temporal order in processing. (They need not be: the various dependencies might correspond to a set of constraints a system attempts to satisfy simultaneously. But see Grainger *et al.*, 2008, *pace* Prinz, 2012, p. 106n, for evidence that hierarchical dependencies in letter perception correspond to processing stages.) The problems posed by a very large space of possible global solutions are often addressed in computer science—including computational work on vision—by 'greedy' algorithms. These algorithms sequentially seek optimal local solutions in the hope that, although features once assigned locally aren't reconsidered, local solutions will combine to approximate a globally optimal solution (Cormen *et al.*, 2001). (A simple, albeit non-visual example is the following algorithm for giving change of amount n: hand over a unit of the largest denomination $m < n$ you have and repeat for n-m until done. A visual example is provided presently.) For this class of algorithm, where selective processing begins will influence the output that ultimately emerges. Suppose participants in experiments such as Sperling's employ something like a greedy algorithm to generate hypotheses concerning letters (higher-level features) from lower-level features. Then their performance can be explained, at least in part, if selective processing starts in different places with and without a cue and produces more accurate outcomes in the regions where it begins—in other words, if errors become more likely as processing progresses. For example, just as some algorithms in machine vision first group only local edges into contours (Li *et al.*, 2012), letter recognition might proceed from a local subset of lower-level features, which, once assigned, cannot be reconsidered when the algorithm attempts to reconstruct letters in adjacent regions.

*Intra*-level dependencies also provide an opportunity for the cue to affect what letter representations are formed: whether some letter is represented as present at a particular location depends not only on the represented presence of the features that make up that letter, but also on the other features represented as present in the scene. For example, letter recognition very likely involves expectations concerning the likely co-occurrences of certain letters. By design, however, there are no reliable contingencies in Sperling's experiment. As a result, a greedy recognition algorithm might utilize expectations in a place where they do not apply. This would depend on having at least one letter already represented, so that it can influence the outcome

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

for the next. The effect would grow as more letters are thought identified; the letters likely to follow 'S' in English, for example, are far greater than those likely to follow 'ST'.

It's not necessary, however, that a greedy algorithm be deployed for the cue to affect the formation of letter representations. For example, one way the visual system addresses recognition problems is by assuming that items nearby one another are more likely to be similar (Dawson, 1991, on motion perception; Wolfe *et al.*, 2011, on scene perception). Applied to the Sperling partial-report condition, the idea is that the cue might signal which inferences about lower-level features influence one another. For example, it appears that, in general, observers represent line segments nearby one another as having more similar orientations than they may have in fact (Orhan and Jacobs, 2013).[11] In the full-report condition, inferences about low-level features (e.g. segment orientation) in each letter position may be influenced by similar inferences about the features in the neighbouring positions. But in the partial-report condition, the cue could cause selective processing that leads only to *horizontal* influences (since Sperling's tones are associated with distinct rows). In principle, this can result in different specific letter reports in the two conditions.[12]

In our examples so far, we've mainly had the cue effect a *restriction* in input—by excluding some lower-level representations or excluding the influence of some same-level neighbours. But another possibility is that the cue affects the relative *weight* such features receive (their *exclusion* is then just the special case where that weight is zero). We can develop this further in the course of developing another idea—viz., that the contents of visual memory according to many (e.g. van den Berg *et al.*, 2014) are not individual letters or other feature values, but probability distributions (or densities) over those entities.[13]

---

[11] Orhan and Jacobs (2013) emphasize how the prior expectation of local similarity influences the representations encoded in visual working memory and note that this expectation is in conflict with the actual sampling procedure that experimenters (including Sperling) typically apply, since they typically construct stimuli by drawing objects or features randomly from a uniform distribution.

[12] There are other ways adjacent features can be relevant. We have already mentioned that some sets of letters co-occur more frequently than others. (Sperling, *pace* Block 2011, did not use all 26 letters, but rather eliminated vowels so as to minimize the chance of words occurring. This does not eliminate, however, the possibility that knowledge of transitional probabilities among letters affected perceptual processing.) Also, some letters, taken individually, occur more frequently than others, meaning also that some sets of lower-level letter-related features occur more frequently than others. Regarding our suggestion above that the cue might cause selective *horizontal* processing, it might be objected that a history of horizontal reading might already have this effect. Even so, it's possible that the cue could magnify the effect. But, in any event, there's no reason to think a horizontal effect applies to the recovery of lower-level features, such as those we emphasize above. The similarity constraint in Orhan and Jacobs (2013) uses Euclidean distance.

[13] Such views are often developed in a Bayesian framework, but we needn't assume that perceptual systems are optimal. For an introduction to Bayesian modeling of perception, see Feldman (forthcoming). The Bayesian tide in cognitive science is not without its critics. For a useful

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

In line with recent work within this framework, one might present an observer with a set of randomly oriented line segments. At test, one position formerly occupied by one of the segments is cued, and the observer reports the orientation she remembers for that object (e.g. by turning a dial to orient a test segment to match the memory). Data are analysed under the assumptions that each segment in the display evokes a transduced signal in the mind of the observer, and that the signal is used to infer a posterior distribution over the most likely orientations for each object. During report, it is assumed that observers 'sample' from the relevant distribution—that is, select a specific hypothesis from the space with a probability equivalent to the probability assigned to the hypothesis. The nature of the content reported (discrete) thus differs from that stored in memory (probabilistic). It should be clear that to the extent this framework applies to experiments on line orientations, colours, and shapes (see our next section for examples), it also applies to Sperling's experiments—at the level both of lower-level features and of letters.

It may seem unintuitive that memory and perception traffic in probabilistic representations. But to the extent that perception is inherently uncertain, the question becomes: why *should* perceptual systems commit themselves concerning some distal feature (e.g. that letter *such-and-such* is *there*) rather than maintain probabilistic information until 'discretization' is required, say by some query? Probabilistic representations contain lots of useful information—for instance, concerning confidence (the inverse of variance in the probability distribution), information which subjects have been shown to make use of in behavioural tasks (Bahrami *et al,.* 2010).[14]

This last point is important. That perception is noisy entails just that the *transitions* from stimuli to signal and between states of the perceptual system are probabilistic; thus the same stimulus over repeated encounters may generate different representations in a way that forms a probability distribution. But subjects' use of information concerning confidence can suggest that our perceptual systems are probabilistic not just in their operations, but also in what they represent: our perceptual systems seem to track, encode, and—when possible—exploit the noisiness to which they are susceptible. The most natural alternative to assigning probabilistic representations to perception, given such behavioural evidence, would be to posit only *post*-perceptual mechanisms that track, encode and exploit the noisiness to which perceptual systems are susceptible. However, recent fMRI work has found evidence in the visual cortex of probabilistic representations reflecting sensory uncertainty (van Bergen *et al.,* forthcoming). We discuss further evidence for probabilistic representations when we turn to Lamme.

If we take seriously that memory represents distributions, a problem arises for Block. For much work of this sort concerns the representations in visual short-term

exchange, see Bowers and Davis (2012a), Griffiths *et al.* (2012) in reply, and Bowers and Davis' (2012b) rejoinder. On whether sampling is consistent with optimality, see Eberhardt and Danks (2011) and Vul *et al.* (2014).

[14] Though, to be sure, the representation of the distribution could also be retained after discretization. Cf. Vul *et al.* (2009), who had participants sample twice from a distribution.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

memory. It seems unlikely that Block would claim that iconic memory, which is supposed to *precede* short-term memory, consists of many discrete specific representations, which are then partially transferred into short-term memory and *transformed* into probability distributions. To be sure, if *both* iconic and short-term memory trade in distributions, one might compare their respective capacities for containing them. But the introduction of probabilities, in addition to requiring a significant departure from Block's conception of these stores, provides an additional way to develop an account of both performance limitations and the impact of post-stimulus cues that doesn't advert to the decreasing capacity of successive memory stores.

Regarding performance limitations, we shall argue below, in reinterpreting Lamme's results, that participants' accurate reports on visual memory tasks are limited to three to five items as a result of increasing interference among probabilistic representations as the number of items in the stimulus increases. On such a view, a post-stimulus cue can affect *which* items are accurately reported by rendering some positions less prone to such interference and thus less prone to error than others. For example, on some accounts, the degree of uncertainty in the relevant probability distributions will be reduced when focal attention is directed to the relevant stimuli (Bays and Husain, 2008). Accordingly, to explain how the cue can affect report, one needs simply to assume that it guides focal attention to the relevant positions in memory—as indeed Makovski *et al.* (2008) suggest regarding other post-stimulus cueing experiments (cf. Phillips, 2011b). Such a possibility is consonant with more general discussions of attention within a Bayesian framework, according to which strong signals produced by exogenous attention-grabbers generate an expectation of low noise (cf. Friston, 2009; Feldman and Friston, 2010 and Hohwy, 2012). Applied to our earlier examples of representational dependence, the suggestion is that the cue can affect the probability distribution by in effect differentially weighting the lower-level inputs to the Bayesian calculation, tightening the resulting curves for attended items and flattening it for others. This is a Bayesian formulation of the familiar idea that attention can differentially affect the gain of elements in a processing system (though see Orhan and Ma, 2015, on stimulus selection versus gain increase).

We can sum up our response to Block's interpretation of Sperling's results by framing it in more general information-theoretic terms. Stimulus displays contain messages, meant to be encoded, stored, and then reported by a receiver. But, first, the messages must be inferred from a signal, the stimulation propagated from the retina. We suggest that cueing should influence inferences regarding messages simply by identifying parts of a signal as noise with respect to a given trial. In contrast, on both Block's view and the views of his previous critics, cueing has its impact by selecting some among several already-inferred messages for transfer into later capacity limited storage. But there are good reasons to doubt that performance limitations reflect the capacity limits of memory stores rather than the exigencies of recovering confident representations from noisy signals in the first place. Our proposal better fits current views of vision and memory. Examining other post-stimulus cueing results on which Block relies provides further support for this claim.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

## 4. Lamme and Colleagues and the Calculation of Capacity

We turn now to more recent post-stimulus cueing experiments by Victor Lamme and colleagues, on which Block leans heavily.

Lamme's group has used several tasks, inspired by Sperling's and by work on change detection and working memory (e.g. Luck and Vogel, 1997). A typical task proceeds as follows (cf. Vandenbroucke *et al.*, 2011). Participants are presented a number of rectangles, each oriented horizontally, vertically, or at a 45° or 135° slant, and arrayed on the perimeter of a circle with a fixation point at its centre. Some time after offset, participants are presented a second display and must report whether one of the rectangles has changed its orientation or not. The main manipulation in these studies involves the timing of a featurally uninformative cue that identifies the position of the object about which observers will need to make the change/no change judgement. In the 'early cue' condition, the cue comes 10 ms after offset of the first display; in the 'retro-cue' condition, it comes between the two displays (for example, 1000 ms after offset of the first—thus outside the temporal window of postdiction—departing 500 ms before onset of the second); and in the 'post-cue' condition, it comes along with the second display. The main finding is that participants perform significantly better in the early than in the retro-cue condition and significantly better in the retro- than in the post-cue condition.

Lamme's group uses these results to argue for a further stage in visual memory—fragile visual short-term memory—intermediate between the two proposed by Sperling. For current purposes, what's crucial is that their experiments share with Sperling's a reliance on the effects of post-stimulus cueing to suggest that more items are present in the contents of early memory than later memory—though our reinterpretation also undermines their basis for hypothesizing the intermediate store. Some of the issues raised by their work differ, however, because of the ways their tasks differ from Sperling's.

In particular, unlike Sperling, Lamme and colleagues use a change/no change report. As a result, the calculation of memory capacity from participants' performance is less transparent. With Sperling's experiments, one basically sums up performance for each row in the 'partial report' condition to determine how many letters were in fact seen altogether. Of course, we challenge the assumption behind this calculation; but, because performance in Sperling's experiments consists of reporting letters seen, it's clear how, *given* the assumption, the calculation should proceed. Things are not so straightforward when the performance consists of change/no change reports. How, from a collection of such reports, can one calculate capacity?

Lamme and colleagues rely on a formula, proposed by Cowan (2001), that relates capacity—Cowan's $K$—to the accuracy of participants' change/no change reports and the task's memory set size (the number of stimuli presented in the task). According to this formula, capacity equals hit rate (i.e. correct change report rate) plus correct rejection rate minus 1, all multiplied by memory set size: $K = (h + r - 1) \times n$. Where does this come from, and why should it provide a guide to capacity? Making

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

explicit the model that underwrites this calculation shows that it assumes what current probabilistic models of visual short-term memory challenge. Evidence for the latter models thus undermines this argument for overflow.

The reasoning behind Cowan's equation is as follows. If a participant has a hit, it's either because that item was among those stored in memory or because, although the item was not among those stored, the participant correctly guessed that it changed. So the hit rate corresponds to the percentage of stimulus items that enter memory summed with the percentage that don't as weighted by a guessing factor *g* (the probability the participant will guess 'change' when the item is not among those stored in memory): $h = K/n + [(n\text{-}K)/n]g$. According to parallel reasoning, the correct rejection rate $r = K/n + [n\text{-}K/n](1\text{-}g)$. Simple algebraic manipulation then allows one to solve for *K*. *g* drops out, leaving the formula for capacity given above (Cowan 2001, p. 166).

Using this formula, Cowan argues—based on a variety of empirical results—for a working memory capacity of about four items, a drop from Miller's (1956) earlier proposal of a 'magic number' of seven items. This is more-or-less the capacity Lamme and colleagues find in the post-cue condition, when the cue is presented along with the second display. It's presumed that the post-cue taps into working memory. But they calculate a capacity of about six items in the retro-cue condition, when the cue occurs between the two displays. Hence their hypothesis of an intermediate memory store with a capacity likewise intermediate between that of iconic and working memory.

The reasoning that underwrites Cowan's formula, however, makes several crucial assumptions. First, the formula for *K* is arrived at by assuming that the hit rate and rejection rate reflect how many items make it into memory. In particular, it assumes a fixed number of slots, and thus a fixed capacity, in visual short-term memory—at least relative to the stimulus set at issue (see footnote 15 below on stimulus complexity). Second, it assumes that these slots are all-or-nothing: an item makes it in or doesn't; there are no probabilities involved in the memory representations. Third, it's assumed that, for items that enter memory, performance is ideal: if you remember it, you will correctly report whether there was a change.

These assumptions are all controversial. Indeed, they are challenged by recent work that, in line with our alternative, suggests a rather different conception of visual working memory (reviewed in Ma *et al.*, 2014 and Bays, 2015). For example, Bays and Husain (2008) found that recall precision declines continuously on a delayed estimation task as the number of items increases. Rather than have participants report a discrete feature of memory items (as with Sperling) or report on whether there was a change (as with Lamme and colleagues), Bays and Husain had participants after stimulus offset indicate on a colour-wheel the colour of an item post-stimulus cued from among several presented. The more items that had been presented, the more variability there was in participants' responses. On the model that underwrites Cowan's reasoning, it is unclear why a larger *n* would have this effect on the variance of participants' responses. If items simply make it into memory or not, why should the fact of further items in the stimulus affect the variability of participants' responses

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

for those items that were indeed remembered? Further, Keshvari *et al.* (2013), using stimuli otherwise similar to Lamme and colleagues', found that the proportion of 'change' reports varies with the *magnitude* of the change. As a result, $K$—that is, capacity, according to Cowan—varies with this change. (In fact, they found it to equal zero for some magnitudes!) But this is odd indeed: how, on the standard view, could varying something in the *later* display affect how many items were *originally* encoded into working memory prior to it?[15]

Whereas the standard fixed-slot model of working memory that underwrites Cowan's formula has difficulty accommodating such results, more recent probabilistic models handle them nicely. According to these models, working memory deploys a continuous limited resource that is shared out among probabilistic representations of variable precision, with a decreasing amount of resource per item as the number of items increases. Rather than representations such as $<red, there_i>$, working memory would contain a representation of a probability distribution or density centered around a point indicating red. In effect, it would contain representations, regarding a particular item in the stimulus, such as $<red, there_i, .7>$, $<blue, there_i, .1>$, and so on for all the colour possibilities regarding that item, with the number indicating the subjective probability of things being so—similarly, regarding other possible items in the stimulus.[16] However, working memory, on these models, deploys a limited resource: as more items are included in the stimulus, the variability of the representations increases (i.e. the precision decreases); the probability curves flatten and spread out, even if centered on the same points as before. The limits in performance that motivate fixed-slot models are thus explained at least in part in terms of interference among probabilistic representations: the more items, the more probability distributions; the more distributions, the less precision; the less precision, the more errors when a distribution is sampled for report. If we speak of capacity limits in memory (as opposed to in performance), they will concern the *precision* of representations, not the *number* of items that can be represented. Concerning the latter, no limitation may be assumed at all! The probabilistic features of such models smoothly accommodate the results from the previous paragraph that were problematic for standard slot models.[17]

Further, they accommodate results that diverge from the third assumption behind Cowan's formula: the assumption that, regarding remembered items, participants judge 'change/no change' correctly. Of course, performance is rarely perfect;

---

[15] This contrasts with apparent variation in capacity owing to differential complexity of the *stimuli*, a difference obviously *prior* to encoding. Cf. Alvarez and Cavanagh (2004), noted by Block (2011, p. 570). But see Awh *et al.* (2007) for a different explanation of these results.

[16] Alternatively, the subjective probability might be an aspect of how the feature is represented—that is built into the 'attitude' as opposed to the content towards which one has the attitude. Nothing in our discussion turns on this.

[17] For further discussion, including of other non-capacity limited models as well as various hybrids, see Ma *et al.* (2014). For an opposing view, see Luck and Vogel (2013).

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

there are always various sources of noise and interference. Indeed, for this reason, Cowan himself, along with colleagues, has recently looked for ways to drop this idealization (Rouder *et al.*, 2011). But of special interest to us are certain potential errors that participants in Lamme and colleagues' experiments are particularly susceptible to owing to the specifics of the task. In arriving at their change/no change judgements, participants must compare a cued probe item to one retained in memory. They therefore must match the cue to the *correct* item both in the first and in the second display. A failure to do so is a *correspondence* error.[18] Lamme and colleagues in effect make the assumption that participants always correctly place the cue in correspondence with both the relevant memory item and the relevant test item. But recent work has shown that correspondence errors are prevalent in related tasks (Bays *et al.*, 2009; Emrich and Ferber, 2012; Bae and Flombaum, 2013). In fact, correspondence errors explain patterns in participants' inaccuracies that are hard to explain on Lamme's assumptions—viz., that that the contents of inaccurate reports tend to cluster around the properties of neighbouring items. The prevalence of correspondence errors is what one would expect on probabilistic models: if representations of features and positions are corrupted by noise and uncertainty, an observer will not necessarily know, when a cue points to an empty space, exactly which memory item previously occupied that space, nor will she know afterwards exactly where it pointed. Indeed, studies with displays and cues nearly identical to those used by Lamme and colleagues—though with a different agenda—have specifically demonstrated that cue interpretation in these contexts is noisy (Vul and Rich, 2010).[19]

In sum, all three assumptions behind the calculation at the heart of Lamme and colleagues' argument, and thus of Block's, are at odds with recent experimental work and challenged by an opposing conception of visual working memory that better fits the data.

What then *should* one say about participants' performance on these tasks? Why does the timing of the probe affect the hit and rejection rates as it does, so that performance deteriorates across the three conditions? A proponent of the capacity-unlimited continuous resource conception may point to several effects of differential cue-timing. First, the timing of the cue can affect the rate and kind of correspondence errors, with the incidence of relevant errors varying with cue-conditions. Second, the timing of the cue can affect the rate of false alarms (e.g. judging 'change' when there is none) by differentially enabling one to ignore, or assign less weight to, 'change' impulses coming from elsewhere in the array. In general, change/no change tasks like Lamme's pose a unique challenge to observers

[18] Of course, sometimes participants who make a correspondence error will nonetheless provide a fortuitously correct report.

[19] The possibility of correspondence errors arises in Sperling's experiment as well: participants must correctly match the grid location where they record a letter report with the remembered location of the letter.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

in that some non-zero probability of a change should be assigned to all positions in a display in all trials. Thus reporting change or not depends on an aggregate evaluation of these probabilities; one might, for example, produce a false alarm not because any single position generates a signal of a greater than 50% chance of change, but merely because several positions summed imply a good chance of a change having taken place in that display. As we suggested with respect to Sperling, a cue can thus have an effect by essentially identifying parts of a display as noise. The early cue would have the greatest effect by significantly reducing the scope of relevant inputs—and thus the computational complexity—involved in producing a task-relevant response, followed by the retro- and then the post-cues. Finally, by redeploying attentional resources, the retro-cue can selectively protect corresponding representations of items in the first display from degradation and interference (Makovski *et al.*, 2008; Pertzov *et al.*, 2013). The post-cue comes too late to have this effect.[20] Thus, it seems possible to explain participants' performance—and the effects of cue-timing thereupon—without having to posit multiple memory stores with declining capacity. Indeed, it is preferable to do insofar as the alternative conception of working memory better accounts for various results.

## 5. Bronfman *et al.* on Ensemble Properties

The final experimental support for overflow that we shall discuss comes from recent work by Bronfman *et al.* (2014). The interest of Sperling and Lamme's post-stimulus cueing paradigms—at least for the debate on which this paper focuses—stems from their promise to indirectly indicate the capacity of early memory stores whose contents are not directly available for report. Bronfman *et al.* offer an alternative indirect probe: participants' access to 'gist' or ensemble properties.

In their crucial experiment, participants were briefly presented with a 4x6 Sperling-like array of *variously coloured* letters, with one row cued *prior* to the presentation and the location of one letter in that row cued *after* offset of the letter array. Participants were told beforehand that they would first be asked to report the post-stimulus cued letter in the pre-stimulus cued row, and then they would be asked to report on colour diversity—in one condition, of the cued row; in the other, of the uncued rows. On trials with high colour diversity, colours were drawn from 19 colours that formed a colour wheel.[21] On trials with low colour

[20] Incidentally, this could explain Vandenbroucke *et al.*'s (2011) apparent dissociation of fragile visual short-term memory from visual working memory by attentional modulation. They found that decreasing attention during the stimulus display greatly reduces the capacity of visual working memory but not that of fragile visual short-term memory. But the performance from which they extract this conclusion is to be expected if what's supposed to tap into fragile visual short-term memory—the retro-cue—provides a compensating attentional boost.

[21] It's a question whether participants, in all cases, experienced the 19 colours as 19 distinct colours, as opposed to categorizing them (sometimes) more coarsely—e.g. {slateblue}, {blue}, {royalblue}, and {steelblue} versus just {blueish} (see the Supplementary Materials to Bronfman

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

diversity, colours were drawn from six adjacent colours. The relevant result was that participants' letter reports yielded a calculated average of three letters remembered in the pre-cued row, in accord with standard views of visual working memory capacity, but, *in addition*, participants' colour diversity reports displayed a sensitivity to the colour diversity of the uncued rows. Bronfman *et al.*—and Block (2014a) in his discussion of their results—argue that this supports the claim that perceptual consciousness overflows cognitive access. The suggestion is that participants' ability to accurately report colour diversity reflects the conscious representation of specific colour information that goes beyond the limits of working memory.

A first worry concerns how, if this information does not make it to visual working memory, it becomes available for report. One might think that, even if (much of the) information about individual colours does not make it to working memory, the ensemble information does—and does so with at most a limited effect on available capacity. One might then object that even just the availability of the ensemble information undermines Bronfman *et al.* and Block's own claims concerning the capacity of visual working memory: how can visual working memory get even *that* for free? (Cf. Byrne *et al.* 2007 on subjects seeming to see all the letters in Sperling's arrays.) Bronfman *et al.* (2014) address this by tentatively suggesting that the colour diversity report is mediated by a distinct, non-selective visual pathway *a la* Wolfe *et al.* (2011). Block (2014a), whose piece is just a brief spotlight report, simply says that the information enters the global workspace for free without taking up how it gets there. This is an issue that must be compellingly addressed if Bronfman *et al.*'s results are to count in Block's favour. We note that there exist probabilistic models of visual working memory that incorporate memory of gist properties (Brady and Tenenbaum, 2013; Orhan and Jacobs, 2013).

But let's bracket this. It remains the case that to find support for the overflow hypothesis in Bronfman *et al.*'s results, one must commit oneself to at least one of various further claims. We start with a strong claim and work down through successively less committed alternatives, in each case arguing that the results fall short of supporting the claim and thus fall short of supporting the overflow hypothesis.

The strong claim is that participants consciously see all, or almost all, of the letters as being the specific letter and colour that they are—that is, with letter and colour features bound together in one object. Neither Bronfman *et al.* nor Block explicitly make this claim. But it is worth mentioning, not only to distinguish it from the less committed alternatives below, but because it would seem in line with the conclusions Block draws from the other experimental results we have discussed. Note, in particular, that Block (2011, pp. 571–3) underscores the importance of the fact that participants in one variant of the paradigm from Lamme's group were able to report change or no-change based on either size or orientation (Landman *et al.*,

*et al.* (2014) for the colours used)—and perhaps also a question whether this might affect the interpretation of their results. On the bias towards generic colours in visual working memory, see Bae *et al.* (2015).

© *2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

2003), indicating that the representation was image-like. Likewise, one might expect Block to suggest that it is at least relevant whether participants in Bronfman *et al.*'s experiment had conscious perceptual representations that bound together specific letters and colours. If this *is* something that matters, then it is important to note that Bronfman *et al.*'s results do not provide sufficient support for the strong claim. It is consistent with participants' performance that they represented the visible presence of all the various individual colours as well as whatever information enabled them to report the cued letter—but without the colour information being bound to representations of specific letters.

Let us turn, then, to this slightly weaker claim: participants represent whatever information results in their performance in reporting letters and, in addition, they consciously see all, or almost all, of the specific colours, albeit not necessarily as bound to specific letters. Bronfman *et al.* seem to endorse this claim. They write: '… the availability of color diversity is best explained as resulting from the fleeting experience of the underlying individual colors' (Bronfman *et al.*, 2014, p. 1395). In correspondence, however, they indicate that it's not their position that a fleeting experience of *all*, or almost all, the colours is required. (Block's (2014a, p. 446) explicit wording is more circumspect: '… there must have been conscious awareness of specific colors beyond the limits of the global workspace.') In any event, again, the experimental results do not provide sufficient support for this claim. In principle, if participants consciously saw just three of the colours in uncued rows, a gist judgement based on this sample could be accurate over 80% of the time.[22]

But still this leaves the even weaker claim that participants consciously saw at least three colours in uncued rows in addition to the three *letters* Bronfman *et al.* calculate that participants consciously saw: in other words, their gist performance, combined with their letter performance, reveals that perceptual consciousness has a capacity at least 50% greater than that of visual working memory—roughly six items versus roughly four items. However, it is a question whether—and, if so, how—distinct features-types (here, colours and letters) compete with each other for working memory resources (e.g. Fougnie and Alvarez, 2011); moreover, visual working memory is believed to have higher capacity for colours than for shapes (Alvarez and Cavanagh, 2004). Unfortunately, Bronfman *et al.* did not also ask participants to report specific

[22] This is an idealized estimate supposing one judges diversity to be low if all three samples are within six of one another but high otherwise. (Barry Smith points out to us that, on some trials, just one uncued colour would suffice, if one could compare colours from the cued row and from uncued rows: suppose the cued row was low diversity and the one uncued colour was not from that group of six. This requires that, when the cued and the uncued rows were both low diversity, all colours were drawn from the same group of six adjacent colours—as indeed was the case.) Our point here is not to suggest that the visual system in fact uses such a rule, but just to shift the burden onto someone who would claim that participants experience all, or almost all, of the specific colours. Note that Bronfman *et al.*'s experiment 3, in which colours in the low diversity condition were not drawn from a set of adjacent colours, would require a different strategy.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

colours seen in order to explore whether any from uncued rows reached working memory. Further, the claim that the experimental results show that perceptual consciousness has a capacity at least 50% greater than that of visual working memory assumes a standard conception of working memory capacity as limited to four items—just what we have suggested cannot be assumed.

Indeed, recall that Bronfman *et al.* did not have participants report what letters they saw in the cued row, the way Sperling did. Rather, they had participants report just the one letter of the pre-cued row that occupied the post-cued location. They then calculated memory capacity on the basis of these reports. So, we may again ask what assumptions their calculation makes. Bronfman *et al.* assume that M of the N stimulus items make it into memory (they take N to be 6, the number of letters in a pre-cued row). They assume ideal performance for remembered letters: if the letter makes it into memory, it will be correctly reported. If a letter doesn't make it into memory, they assume participants are equally likely to guess any of the nine letters used to construct the arrays. It follows that the expected accuracy rate $A = (MC + N - M)/NC$, where C is the number of response alternatives (here $C = 9$). One can then calculate M by plugging the observed accuracy rate into: $M = N(CA - 1)/(C - 1)$. But this calculation makes assumptions we have already questioned—for example, that representations are all-or-nothing, that performance on remembered items is ideal, and that perceptual systems treat letters as equally likely. All sides can agree about limitations in *performance*. Standard models would *explain* limitations in performance in terms of limited slots in all-or-nothing working memory. The alternative resource conception explains limited performance instead in terms of decreased precision and interference among probabilistic representations. It's not clear why, from the perspective of this alternative, it should be surprising that the ability to report ensemble properties does not interfere with the ability to report other more specific features.[23]

The claim that participants consciously experienced three colours in uncued rows thus doesn't require positing successive memory stores of decreasing capacity. But, in any event, it can be questioned even whether participants *consciously* experienced those colours, as others have responded to Block on Sperling. Must we assume that participants' colour diversity reports reflect conscious experience of specific colours at all? Bronfman *et al.*, sensitive to this challenge, provide experimental reason to think that they do. In one variant of their task, participants were instructed to press an escape button if they had no impression of the uncued letters' colours. No participant pressed the button when the letters were coloured, and participants were generally accurate in pressing on 'catch' trials with colourless uncued letters. In

[23] Bronfman *et al.* might point here to their computational model simulating the introduction of noise, which they argue would affect colour diversity judgements more significantly than judgements of average colour (see their Supplementary Material). Be that as it may, the model still yields sufficiently good results for colour diversity judgements, apparently in line with actual performance.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

another experiment, participants were asked for visibility ratings of more briefly presented post-masked arrays with varying inter-stimulus intervals, and their ratings strongly correlated with relative accuracy of colour diversity judgements.[24]

But there is room to challenge these grounds for attributing consciousness of specific colour features (for related points, see Phillips, forthcoming). Regarding the first experiment: that participants, who expected colour, detected the absence of this generic feature does not entail that they represented (consciously or otherwise) *specific* colours on non-catch trials. Regarding the second: there's perhaps room to wonder whether participants' ratings reflected their confidence in their reports, rather than conscious visibility. Participants might find 'I'm not so sure, but I'd say there was high colour diversity' a more natural or appropriate judgement to communicate than, in effect, 'I didn't see the colours, or not many of them, but I'd say there was high colour diversity.'[25] In any event, even if their ratings do reflect their conscious experience, it's again not obvious that they reflect conscious experience of specific colours. It's possible that their ratings reflect rather their conscious experience of a more generic property—for instance, the experience of (generic) colour. (Recall: Block rejects at least the general version of the Humean principle that experience as of a determinable requires experience as of one of its determinates.) Indeed, it's possible that what drives their performance in *both* experiments is the conscious experience of *gist properties* concerning colour, as opposed to the conscious experience of the various specific colours. This is consistent with the still-developing literature on gist (cf. Wolfe *et al.*, 2011, p. 82, a paper to which Bronfman *et al.* advert). To assume that gist (statistical, ensemble) properties can only be conscious if more specific information from which they are gleaned is conscious would beg the question.[26]

## 6. Conclusion: Probabilistic Working Memory and Perceptual Consciousness

We have argued that the experiments to which Block adverts do not in themselves support his overflow hypothesis. There are alternative interpretations—in line with current views on perception and memory—that explain these experimental results without invoking successive memory stores with declining capacity. On the view we favour, performance limitations on visual memory tasks reflect the challenges that

[24] In another experiment, letters were either blue or red, and, in each trial, one colour was dominant (16 of 24 letters were of the dominant colour). Participants' visibility ratings correlated with their relative accuracy of colour dominance judgements.

[25] Participants gave visibility ratings first and colour diversity estimates second. For visibility, they chose among 'did not see the colours,' 'partially saw the colours,' and 'saw the colours well.'

[26] For evidence that accurate colour diversity judgements about such displays do not require conscious experience of individual colours, see now Ward *et al.* (forthcoming).

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

noisy, ambiguous signals present to the recovery of confident representations concerning multiple objects; and post-stimulus cues have their effects on performance by modulating the generation, maintenance and comparison of these representations.

Now, Block of course does not claim that his conclusions are *entailed* by these experiments; he is offering an inference to the best explanation. So it would be reasonable of him to suggest that our alternative interpretations do not yet displace his own unless they find a place in a larger overall view that provides, all things considered, an even better explanation. We draw from a growing body of work on perception and memory that accommodates the kinds of results that are problematic for the view of visual memory common among Block and his previous critics; to that extent, we do indeed provide a better explanation of the experimental results on which Block relies. But Block can counter that the relevant *explananda* include the facts of perceptual consciousness (the facts themselves, not just the (defeasible) introspections concerning phenomenology that provide partial theoretical access to those facts), about which we have prescinded for the most part from advancing positive claims. Moreover, it's not obvious what, from our perspective, one ought to say. Of course our considerations can't be *dismissed* on this basis, but it might level the field so far as the best explanation was concerned.

So, let's conclude by briefly and somewhat speculatively ruminating on perceptual consciousness. We comment on two questions in particular. First, with which store is phenomenology associated? And, second, *what is it like* to visually represent a probability distribution? Doing so brings to the fore as well some architectural questions that any development of our alternative must address.

Block associates perceptual consciousness with iconic memory, the earlier store; previous critics associate perceptual consciousness with working memory, the later store. What should our view be? An important prior question is whether we—like Block and his critics—should recognise a distinct store prior to working memory in the first place. Throughout, we have challenged the claim of successive stores of declining capacity. One can reject this claim by rejecting only the claim of declining capacity, leaving in place the claim of successive stores. But why do so, rather than just posit transitions from noisy signals to probabilistic representations, without any transition from a first such store to a second? The latter view is simpler and requires fewer resources.

One obstacle to dropping the first store arises from such obvious facts as that, when Sperling's participants record their results, what they see—a blank grid on which they are to write what letters they saw—differs from what they are remembering. New phenomenology, as it were, pushes out the old, while the contents of working memory remain available for sampling and report. This might suggest the existence of two stores after all, with the earlier store, prior to working memory, associated with perceptual consciousness. On this view, by the time Sperling's participants utilize the second store for report, the first store has up-dated to a new visual representation. But now we must ask why we would represent more-or-less the same thing twice. (The 'declining capacity' view of course denies that more-or-less everything is re-represented and typically, as a further claim, assigns successive stores

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

a role in weeding out excess information.) One possibility is that this helps optimize the balance between obtaining new and retaining old information. Another (not exclusive) possibility is that transfer across these stores functions to reformat contents so that they can combine with contents from other sources. The second store might thus serve as an interface between a modular perceptual system, with its proprietary representational format, and domain-general conceptual systems (cf. Carruthers, 2006). It's unclear, however, why this role could not be played by the process of discretizing from working memory itself.

Moreover, it may be that visual temporal dynamics can explain the difference between what one sees and what one remembers without positing two stores. Presumably, it takes time for a probability distribution to be sampled and for that to lead to a report; perception needn't remain frozen in the meanwhile. Thus might Sperling's subjects see a grid (the current content of their one visual store) even as the effects of what they earlier saw (the previous content of that store) propagate towards their report thereof. This might, however, seem strained applied to Lamme's task. For example, on the post-cue trials we discussed, there is a full two seconds between arrays (in some variants, substantially longer), with different phenomenology associated with different periods in that interval (a black screen with a centered fixation point, followed by the cue, and then another black screen)—and yet participants retain sufficient content regarding the first array to perform significantly better than chance in their comparison with the second. The content of the store thus gets replaced *several times*, over the course of a substantial interval, before the subject must draw upon the original content. Perhaps this suggests a need for multiple stores after all. Or perhaps the worry depends on an overly simplistic conception of such a store's capabilities: it's clear that subjects can hold temporal information in visual memory (they remember not only having seen such-and-such, but also having seen this before that (Marshuetz, 2005)), so perhaps one store can retain sufficient information about successive displays to account for their performance on Lamme's task.

Let's turn now to the question of *what is it like* to visually represent a probability distribution. It's tempting to answer: it's like *this*. But visual experience can seem to present us with discrete objects possessing discrete properties, not with various collections of objects and their various properties all associated with varying degrees of subjective confidence. In fact, there's room to question the extent to which this is so. Apparent perceptual indeterminacies—e.g. Ryle's speckled hen and crowding—might be understood as involving probabilistic, as opposed to merely non-specific, phenomenology. And perhaps this approach can be extended further than one might initially have thought. This idea is certainly worth exploring, but we mention an alternative as well.[27]

[27] See Morrison (forthcoming) for an argument that positing probabilistic *perceptual* confidence as an aspect of conscious perception would best explain levels of confidence associated with *beliefs* based upon such perceptions. However, perceptual confidence in Morrison's sense and the probabilities associated with bistable displays seem to dissociate. With such displays, perceivers'

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

A natural thought is to associate phenomenology, not directly with the memory store containing probabilistic representations, but with the discrete representations sampled therefrom. If there are two stores, one might hypothesize either that the discrete representations associated with perceptual consciousness are sampled from the first store or that they are sampled from the second. If they are sampled from the first, this might ease worries about timing raised above, since reports would presumably be based upon distinct representations sampled from the second store. If there is but one store, one might still hypothesize distinct discretizations from this one store: one generating conscious visual representations, the other generating the discrete representations drawn upon in report. Note that, whether drawn from one store or two, if there were multiple discretizations, one would expect discrepancies between them, in accordance with the probabilities associated with sampling—presumably a prediction that could be tested in some manner. But the simplest implementation would have one probabilistic memory store along with one sampling that is both associated with perceptual consciousness and implicated in participants' reports. The phenomenological aspect of the sampled state needn't persist as long as its content, which again would ease worries about timing.[28]

One might suggest that, on such a view, there are successive memory stores after all: first the probabilistic representations and then the discrete representations sampled therefrom. It is not important to us whether the latter should in fact be deemed a memory store. Even if it is, this provides little comfort to Block. Consider first the simpler view where one sampling is both associated with perceptual consciousness and implicated in participants' reports. This obviously denies Block's central claim that perceptual consciousness overflows cognitive access. What's more, though the view agrees with Block's previous critics on this point, it significantly departs from both their views and Block's on other central matters. First, the later store is now one that follows *after* working memory. The successive stores would thus be quite different from those posited in previous discussions. Moreover, no argument has been given that the later store would have lower capacity than working memory. Even if we allow that these capacities can be meaningfully compared, the magic number $4 \pm 1$ was inferred from the number of *accurate* reports, not the number of reports made altogether. Sperling's participants were instructed to fill in *all* the squares of the

visual experiences flip between interpretations, each of which Morrison would associate with high perceptual confidence. But Moreno-Bote *et al.* (2011) show that the durations of such experiences—that is, until a flip—are predicted by sampling from a *stable* probability distribution (which seems to favour the approach we take up next). We should thus distinguish the probabilistic perceptual confidence Morrison hypothesizes and the probabilities encoded in such representations. (Morrison's perceptual confidence would thus not be the same as the inverse of variance—sometimes also given the label 'confidence'—in such representations.) This seems likewise to distinguish his claims from those to which we advert. It may well be that multiple assignments of degrees are needed.

[28] The view can *allow* the probabilistic representation to be sampled more than once, for example in response to an explicit query—cf. Vul *et al.* 2009. The second sampling needn't be associated with perceptual phenomenology, even if it might generate conscious visual imagery.

© *2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

grid, guessing where necessary—and even the guesses might have expressed discrete representations associated with lower confidence (on account of the low precision of the probabilistic representations from which they were sampled). If we go by just the number of reports, then one would have to say Sperling's participants demonstrated a report capacity of at least 12 on such tasks. Thus, even if they are deemed successive stores, further argument would be needed to revive the claim that they constitute successive stores of declining capacity. The less simple view on which there are two samplings that in principle could diverge adds a further twist. For it follows that perceptual consciousness and cognitive access could diverge, contrary to the views of previous critics.[29] It would thus distance our alternative even further from those views and, on this point, move our alternative closer to Block, for whom the possibility of perceptual consciousness without cognitive access is central. But the distance from Block would remain significant as well. Again, the stores differ in character from those of his model, and no claim of declining capacity is made: divergence is not overflow. What's more, Block or others might be wary in any event of the two-samplings variant, fearing that the predicted divergence might be implausibly large or that the position threatens an unattractive epiphenomenalism about phenomenal consciousness.

From the perspective of our alternative proposal, there are thus things to be said and directions to explore further regarding perceptual consciousness. This suffices to dislodge Block's claim to offer the best explanation of post-stimulus cueing results.

*Department of Philosophy*
*Johns Hopkins University*

*Department of Psychological and Brain Sciences*
*Johns Hopkins University*

## References

Alvarez, G. and Cavanagh, P. 2004: The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science*, 15, 106–11.

Awh, E., Barton, B. and Vogel, E. 2007: Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science*, 18, 622–8.

[29] If cognitive access requires that the state's content be in some relevant sense gotten *from* the conscious state, then *none* of perceptual consciousness would be cognitively accessed. Nor would perceptual consciousness even be cognitively access*ible*, contrary to Block. These results don't follow if cognitive access requires just that the states agree in content and perhaps have a common source, and if a perceptual state is cognitively accessible if it could have been sampled.

© 2017 John Wiley & Sons Ltd

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Bae, G. and Flombaum, J. 2013: Two items remembered as precisely as one: how integral features can improve visual working memory. *Psychological Science*, 24, 2038–47.

Bae, G., Olkkonen, M., Allred, S. and Flombaum, J. 2015: Why some colors appear more memorable than others: a model combining categories and particulars in color working memory. *Journal of Experimental Psychology: General*, 144, 744–63.

Bahrami, B., Olsen, K., Latham, P., Roepstorff, A., Rees, G. and Frith, C. 2010: Optimally interacting minds. *Science*, 329, 1081–5.

Bays, P. 2015: Spikes not slots: noise in neural populations limits working memory. *Trends in Cognitive Sciences*, 19, 431–8.

Bays, P., Catalao, R. and Husain, M. 2009: The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9, 7.

Bays, P. and Husain, M. 2008: Dynamic shifts of limited working memory resources in human vision. *Science*, 321, 851–4.

Block, N. 1995: On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18, 227–87.

Block, N. 2005: Two neural correlates of consciousness. *Trends in Cognitive Sciences*, 9, 46–52.

Block, N. 2007a: Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavorial and Brain Sciences*, 30, 481–99.

Block, N. 2007b: Overflow, access, and attention. *Behavorial and Brain Sciences*, 30, 530–48.

Block, N. 2008: Consciousness and cognitive access. *Proceedings of the Aristotelian Society*, 108, 289–317.

Block, N. 2011: Perceptual consciousness overflows cognitive access. *Trends in Cognitive Science*, 12, 567–75.

Block, N. 2013: Seeing and windows of integration. *Thought*, 2, 29–39.

Block, N. 2014a: Rich conscious attention outside focal attention. *Trends in Cognitive Sciences*, 18, 445–7.

Block, N. 2014b: Seeing-as in the light of vision science. *Philosophy and Phenomenological Research*, 89, 560–72.

Block, N. 2015: Consciousness, big science and conceptual clarity. In J. Freeman and G. Marcus (eds), *The Future of the Brain: Essays by the World's Leading Neuroscientists*. Princeton: Princeton University Press.

Bouma, H. 1970: Interaction effects in parafoveal letter recognition. *Nature*, 226, 177–8.

Bowers, J. and Davis, C. 2012a: Bayesian just-so stories in psychology and neuroscience. *Psychological Bulletin*, 138, 389–414.

Bowers, J. and Davis, C. 2012b: Is that what Bayesians believe? Reply to Griffiths, Chater, Norris, and Pouget (2012). *Psychological Bulletin*, 138, 423–6.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Brady, T. and Tenenbaum, J. 2013: A probabilistic model of visual working memory: incorporating higher order regularities into working memory capacity estimates. *Psychological Review*, 120, 85–109.

Bronfman, Z., Brezis, N., Jacobson, H. and Usher, M. 2014: We see more than we can report: 'cost-free' color phenomenality outside focal attention. *Psychological Science*, 25, 1394–403.

Burge, T. 2007: Psychology supports independence of phenomenal consciousness. *Behavioral and Brain Sciences*, 30, 500–1.

Burge, T. 2010: *Origins of Objectivity*. Oxford: Oxford University Press.

Byrne, A., Hilbert, D. and Siegel, S., 2007: Do we see more than we can access? *Behavorial and Brain Sciences*, 30, 501–2.

Carruthers, P. 2006: *The Architecture of the Mind*. Oxford: Oxford University Press.

Choi, H. and Scholl, B. 2006: Perceiving causality after the fact: postdiction in the temporal dynamics of causal perception. *Perception*, 35, 385–99.

Chow, S. 1985: Iconic store and partial report. *Memory and Cognition*, 13, 256–64.

Cohen, M. and Dennett, D. 2011: Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15, 358–64.

Coltheart, M. 1980. Iconic memory and visible persistence. *Perception & Psychophysics*, 27, 183–228.

Cormen, T., Leiserson, C., Rivest, R. and Stein, C. 2001: Greedy algorithms. *Introduction to Algorithms*, 1st ed., 329–55.

Cowan, N. 2001: The magical number 4 in short-term memory: a reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87–114.

Dawson, M. 1991: The how and why of what went where in apparent motion: modeling solutions to the motion correspondence problem. *Psychological Review*, 98, 569.

De Gardelle, V., Sackur, J. and Kouider, S. 2009: Perceptual illusions in brief visual presentations. *Consciousness and Cognition*, 18, 569–77.

Dehaene, S., Changeux, J.-P., Naccache, L., Sackur, J. and Sergent, C. 2006: Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in Cognitive Sciences*, 10, 204–11.

Dretske, F. 2006: Perception without awareness. In T. Gendler and J. Hawthorne (eds), *Perceptual Experience*. Oxford: Oxford University Press.

Duncan, J. 1980: The demonstration of capacity limitation. *Cognitive Psychology*, 12, 75–96.

Eberhardt, F. and Danks, D. 2011: Confirmation in the cognitive sciences: the problematic case of Bayesian models. *Minds and Machine*, 21, 389–410.

Emrich, S. and Ferber, S. 2012: Competition increases binding errors in visual working memory. *Journal of Vision*, 12, 12.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Feldman, J. forthcoming: Bayesian models of perceptual organization. In J. Wagemans (ed.), *The Oxford Handbook of Perceptual Organization*. Oxford: Oxford University Press.

Feldman, H. and Friston, K. 2010: Attention, uncertainty, and free-energy. *Frontiers in Human Neuroscience*, 4, 1–23.

Fougnie, D. and Alvarez, G. 2011: Object features fail independently in visual working memory: evidence for a probabilistic feature-store model. *Journal of Vision*, 11, 1–12.

Freeman, J. and Simoncelli, E. 2011: Metamers of the ventral stream. *Nature Neuroscience*, 14, 1195–204.

Friston, K. 2009: The free-energy principle: a rough guide to the brain? *Trends in Cognitive Sciences*, 13, 293–301.

Geisler, W., Perry, J., Super, B. and Gallogly, D. 2001: Edge co-occurrence in natural images predicts contour grouping performance. *Vision Research*, 41, 711–24.

Girshick, A., Landy, M. and Simoncelli, E. 2011: Cardinal rules: visual orientation perception reflects knowledge of environmental statistics. *Nature Neuroscience*, 14, 926–32.

Grainger, J., Rey, A. and Dufau, S. 2008: Letter perception: from pixels to pandemonium. *Trends in Cognitive Sciences,* 12, 381–7.

Griffiths, T., Chater, N., Norris, D. and Pouget, A. 2012: How the Bayesians got their beliefs (and what those beliefs actually are): comment on Bowers and Davis (2012). *Psychological Bulletin*, 138, 415–22.

Grush, R. 2007: A plug for generic phenomenology. *Behavioral and Brain Sciences*, 30, 504–5.

Hohwy, J. 2012: Attention and conscious perception in the hypothesis testing brain. *Frontiers in Psychology*, 3, 1–14.

Hume, D. 1740/1978. *A Treatise of Human Nature*. R. Selby-Bigge and P. Nidditch (eds). Oxford: Oxford University Press.

Keshvari, S., van den Berg, R. and Ma, W. 2013: No evidence for an item limit in change detection. *PLoS computational biology*, 9, e1002927.

Kouider, S., de Gardelle., V. and Dupoux, E. 2007: Partial awareness and the illusion of phenomenal consciousness. *Behavioral and Brain Sciences*, 30, 510–1.

Kouider, S., de Gardelle, V., Sackur, J. and Dupoux, E. 2010: How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences*, 14, 301–7.

Kouider, S., de Gardelle, V. and Sackur, J. 2012: Do we still need phenomenal consciousness? *Trends in Cognitive Sciences*, 16, 140–1.

Kouider, S. and Dehaene, S. 2007: Levels of processing during non-conscious perception: a critical review of visual masking. *Philosophical Transactions of the Royal Society B*, 362, 857–75.

Landman, R., Spekreijse, H. and Lamme, V. 2003: Large capacity storage of integrated objects before change blindness. *Vision Research*, 43, 149–64.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Li, Y., Sawada, T., Latecki, L., Steinman, R. and Pizlo, Z. 2012: A tutorial explaining a machine vision model that emulates human performance when it recovers natural 3D scenes from 2D images. *Journal of Mathematical Psychology*, 56, 217–31.

Luck, S. and Vogel, E. 1997: The capacity of visual working memory for features and conjunctions. *Nature*, 309, 279–81.

Luck, S. and Vogel, E. 2013: Visual working memory capacity: from psychophysics and neurobiology to individual differences. *Trends in Cognitive Sciences*, 17, 391–400.

Ma, W., Husain, M. and Bays, P. 2014: Changing concepts of working memory. *Nature Neuroscience*, 17, 347–56.

Ma, Z. and Flombaum, J. 2013: Off to a bad start: uncertainty about the presence of targets at the onset of multiple object tracking. *Journal of Experimental Psychology: Human Perception & Performance*, 39, 1421–32.

Makovski, T., Sussman, R., and Jiang, Y. 2008: Orienting attention in visual working memory reduces interference from memory probes. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 34, 369–80.

Marshuetz, C. 2005: Order information in working memory: an integrative review of evidence from brain and behavior. *Psychological Bulletin*, 131, 323–39.

McClelland, J. and Rumelhart, D. 1981: An interactive activation model of context effects in letter perception: Part 1. An account of basic findings. *Psychological Review,* 88*,* 375–407.

Miller, G. 1956: The magical number seven, plus or minus two: some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.

Moreno-Bote, R., Knill, D. and Pouget, A. 2011: Bayesian sampling in visual perception. *Proceedings of the National Academy of Sciences*, 108, 12491–6.

Morrison, J. forthcoming: Perceptual confidence. *Analytic Philosophy.*

Navon, D. 1984: Resources—a theoretical soup stone? *Psychological Review*, 91, 216.

Neisser, U. 1967: *Cognitive Psychology*. Englewood Cliffs, NJ: Prentice-Hall.

Norris, D. and Kinoshita, S. 2012: Reading through a noisy channel: why there is nothing special about the perception of orthography. *Psychological Review*, 119, 517–45.

Orhan, A. and Jacobs, R. 2013: A probabilistic clustering theory of the organization of visual short-term memory. *Psychological Review*, 120, 297–328.

Orhan, A. and Ma, W. 2015: Neural population coding of multiple stimuli. *The Journal of Neuroscience*, 35, 3825–41.

Papineau, D. 2007: Reuniting (scene) phenomenology with (scene) access. *Behavioral and Brain Sciences*, 30, 520.

Pertzov, Y., Bays, P., Joseph, S. and Husain, M. 2013: Rapid forgetting prevented by retrospective attention cues. *Journal of Experimental Psychology: Human Perception and Performance*, 39, 1234–41.

Phillips, I. 2011a: Perception and iconic memory: what Sperling doesn't show. *Mind & Language*, 26, 281–311.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Phillips, I. 2011b: Attention and iconic memory. In C. Mole, D. Smithies and W. Wu (eds), *Attention: Philosophical and Psychological Essays*. Oxford: Oxford University Press.

Phillips, I. forthcoming: No watershed for overflow: recent work on the richness of consciousness. *Philosophical Psychology.*

Prinz, J. 2012: *The Conscious Brain*. Oxford: Oxford University Press.

Rao, R., Eagleman, D. and Sejnowski, T. 2001: Optical smoothing in visual motion perception. *Neural Computation*, 13, 1243–53.

Rouder, J., Morey, R., Morey, C. and Cowan, N. 2011: How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin & Review*, 18, 324–30.

Rumelhart, D. and McClelland, J. 1982: An interactive activation model of context effects in letter perception: Part 2. The context enhancement effect and some tests and extensions of the model. *Psychological Review, 89,* 60–94.

Sekuler, R., Sekuler, A. and Lau, R. 1997: Sound alters visual motion perception. *Nature*, 385, 308.

Sergent, C., Ruff, C., Barbot, A., Driver, J. and Rees, G. 2011: Top-down modulation of human early visual cortex after stimulus offset supports successful postcued report. *Journal of Cognitive Neuroscience*, 23, 1921–34.

Sergent, C., Wyart, V., Babo-Rebolo, M., Cohen, L., Naccache, L. and Tallon-Baudry, C. 2013: Cueing attention after the stimulus is gone can retrospectively trigger conscious attention. *Current Biology*, 23, 150–5.

Shore, D., Spence, C., and Klein, R. 2001: Visual prior entry. *Psychological Science*, 12, 205–12.

Simons, D. and Levin, D. 1997: Change blindness. *Trends in Cognitive Sciences*, 1, 261–7.

Simons, D. and Rensink, R. 2005: Change blindness: past, present, and future. *Trends in Cognitive Sciences*, 9, 16–20.

Sperling, G. 1960: The information available in brief visual presentations. *Psychological Monographs*, 74.

Stazicker, J. 2011: Attention, visual consciousness and indeterminacy. *Mind & Language*, 26, 156–84.

Stocker, A. and Simoncelli, E. 2006: Noise characteristics and prior expectations in human visual speed perception. *Nature Neuroscience*, 9, 578–85.

Titchener, E. 1908: *Lectures on the Elementary Psychology of Feeling*. New York: MacMillan.

Tsal, Y. and Benoni, H. 2010: Diluting the burden of load: perceptual load effects are simply dilution effects. *Journal of Experimental Psychology: Human Perception and Performance*, 36, 1645.

Tye, M. 2006: Content, richness, and fineness of grain. In T. Gendler and J. Hawthorne (eds), *Perceptual Experience*. Oxford: Oxford University Press.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License

Ullman, S. 1979: The interpretation of structure from motion. *Proceedings of the Royal Society of London B: Biological Sciences*, 203, 405–26.

van Bergen, R., Ma, W., Pratte, M. and Jehee, J. forthcoming: Sensory uncertainty decoded from visual cortex predicts behavior. *Nature Neuroscience.*

van den Berg, R., Awh, E. and Ma, W. 2014: Factorial comparison of working memory models. *Psychological Review*, 121, 124–49.

Vandenbroucke, A., Sligte, I. and Lamme, V. 2011: Manipulations of attention dissociate fragile visual short-term memory from visual working memory. *Neuropsychologia*, 49, 1559–68.

Vul, E., Hanus, D. and Kanwisher, N. 2009: Attention as inference: selection is probabilistic; responses are all-or-none samples. *Journals of Experimental Psychology: General*, 138, 546–60.

Vul, E., and Rich, A. 2010: Independent sampling of features enables conscious perception of bound objects. *Psychological Science*, 21, 1168–75.

Vul, E., Goodman, N., Griffiths, T. and Tenenbaum, J. 2014: One and done? Optimal decisions from very few samples. *Cognitive Science*, 38, 599–637.

Ward, E., Bear, A. and Scholl, B. forthcoming: Can you perceive ensembles without perceiving individuals?: The role of statistical perception in determining whether awareness overflows access. *Cognition*.

Wolfe, J., Vo, M., Evans, K. and Greene, M. 2011. Visual search in scenes involves selective and non-selective pathways. *Trends in Cognitive Sciences*, 15, 77–84.

Yuille, A. and Kersten, D. 2006: Vision as Bayesian inference: analysis by synthesis? *Trends in Cognitive Sciences*, 10, 301–8.

*© 2017 John Wiley & Sons Ltd*

14680017, 2017, 3, Downloaded from https://onlinelibrary.wiley.com/doi/10.1111/mila.12144 by Johns Hopkins University, Wiley Online Library on [05/11/2026]. See the Terms and Conditions (https://onlinelibrary.wiley.com/terms-and-conditions) on Wiley Online Library for rules of use; OA articles are governed by the applicable Creative Commons License