What have we learned from two decades of object-substitution masking? Time to update: object individuation prevails over substitution

Stephanie C. Goodhew

Research School of Psychology, The Australian National University

Word count: (main text): 10,520

Corresponding Author: Stephanie C. Goodhew

Address: Research School of Psychology (Building 39)

The Australian National University, Canberra, 2601

Email: stephanie.goodhew@anu.edu.au

Running head: Two decades of OSM: Time to update

Abstract

Object-substitution masking (OSM) refers to when the delayed disappearance of a sparse mask that spatially-surrounds but does not overlap the target impairs target perception. Two major theoretical accounts have been offered to explain OSM: the object-substitution account, which stipulates that masking occurs when a separate mask representation replaces the target, and the object-updating account, which espouses that masking is the product of a single representation initially containing information about the target that is modified to reflect the mask. Here I critically review the evidence that has accumulated over two decades for the two models, and find the evidence overwhelmingly in favour of the object-updating account. This object-updating account places OSM in the larger framework of related phenomena such as a repetition blindness, apparent motion, and object correspondence through occlusion which gauge how the visual system assigns episodic object representations in the face of dynamic and ambiguous input. Implications for visual cognition more broadly are discussed.

Keywords: object-substitution masking; four-dot masking; visual masking; object perception; dynamic vision; object individuation; object substitution; object updating; object correspondence; apparent motion.

Translational Abstract

Two decades ago scientists discovered that when an object appears briefly, whether or not people see it can depend on what happens with another, secondary object in the visual scene. When this secondary object appears at the same time as the main object and remains visible after the main object has disappeared, then people typically fail to see the main object at all. There has been debate among scientists about what mechanism is responsible for this effect. This commentary critically reviews the evidence and shows that this effect is a consequence of the brain mechanisms that allow for tracking of object representations as they move and change. In other words, when the secondary object remains visible, the brain infers that it is an altered continuation of the main object, and so the initial representation of the main object is updated to instead reflect the secondary object, which is consciously experienced.

Object-substitution masking (OSM) is approaching its twentieth birthday. The phenomenon was first described and named in 1997 (Enns & Di Lollo, 1997), and its properties were subsequently more extensively tested after that (Di Lollo, Enns, & Rensink, 2000; Enns & Di Lollo, 2000). The purpose of this review, therefore, is to highlight the advances that have been made in our knowledge of the processes that underlie OSM, and therefore the insight that it has and can continue to provide into the mechanisms of human vision. I will also discuss the relationship between OSM and other related phenomena (e.g., repetition blindness, flash-lag effect, attentional blink, etc), with a view to highlighting where there are common processes and where there are not, and what this tells us about visual cognition more broadly. Other interesting questions, such as the debate regarding the level of implicit processing that occurs in OSM, have been extensively reviewed elsewhere (Goodhew, Pratt, Dux, & Ferber, 2013), and therefore will not be revisited here. Instead, I will focus on the understanding of why OSM occurs, analyse what it reflects, and illustrate how this informs us. In particular, I will critically review the existing evidence for the two major theoretical accounts that have been offered to explain OSM: the object-substitution and object-updating accounts. In doing so, I will highlight their similarities and differences, and reveal how there is overwhelming support for object-individuation processes in OSM. I will then discuss the implications of this analysis for OSM in the future.

I would also like to explain my motivation for this review. Most papers on OSM are not aimed at differentiating the two accounts, and therefore typically give at least equal weighting to both views, if not outright preference to object-substitution given its historical precedence. My concern is that may have distorted the consensus in the literature, and could give a reader a misleading take on their relative merits. As a concrete example of the pitfalls of such a bias, Bouvier and Treisman (2010) relied on an object substitution framework, and concluded that single-feature targets (e.g., a single oriented line) where not susceptible to OSM, because they were not sufficiently complex to require re-entrant processing and thus were impervious to masking. Instead they concluded that a target consisting of a conjunction of features was a necessary condition for OSM (Bouvier & Treisman, 2010). Re-examining this conclusion through an object-updating lens, however, we found that even simple single-feature targets can indeed be masked, and that *similarity* between the target and mask features was the determining factor, rather than absolute target complexity (Goodhew, Edwards, Boal, & Bell, 2015). Similarity between the target and mask makes perfect sense as a key consideration within an object-updating framework, but is not an important consideration within an object-substitution framework. This is an example of where failing to consider object-updating processes in OSM led to the incorrect theoretical conclusion.

In a similar vein, as cognitive psychologists, one of our goals ought to be to provide theoretical frameworks and experimental tools that can be used to provide an enhanced understanding of how cognitive processes go awry in psychopathology, with a view to ultimately improving treatment outcomes. On this note, it is concerning that the pervasive over-emphasis on the object substitution account in the literature has led to studies in the clinical literature using what is purportedly OSM to understand "re-entrant activations" in Schizophrenia, where the four-dot mask appears *after* the target (Green, Wynn, Breitmeyer, Mathis, & Nuechterlein, 2011; Jahshan et al., 2014; Lee et al., 2014). Such an experimental set-up may be non-problematic from a substitution perspective. But it confounds the processes of OSM with metacontrast masking (see Breitmeyer & Ögmen, 2006) due to the mask onset occurring after target offset. Whereas an object-updating perspective, which emphasises the relationship between the target and mask, would instead advocate for the presence of the four-dots during target exposure. These are just a few examples of the issues that stem from using the inappropriate theoretical lens. It is my aim, therefore, to correct this bias, and to offer a single reservoir of analysis that provides clear support for the object-updating account. This is timely, and will allow us to move into the third decade of OSM research with a clear consensus about the nature of the tool we use to research visual cognition.

The structure of this review will take the following form. I will firstly review what OSM is. I will then describe the major theoretical accounts, the evidence for these respective views, and offer critical analysis to arrive at a unifying theoretical framework for OSM: object-updating. While this strong support for object-updating does not necessarily preclude the additional presence of substitution mechanisms in OSM as well, I will then review how OSM relates to other phenomena, including visual attention, repetition blindness, the flash-lag effect, object correspondence across saccades and through occlusion, and motion-induced blindness and explain how this also provides converging evidence for this single, unitary theoretical understanding of OSM, and the mechanisms of dynamic object representation processes.

What is object-substitution masking?

In the nomenclature of visual masking, the stimulus that participants are to identify or detect is called the *target*, and the other stimulus presented close to the target in space and time that distorts perception of the target is called the *mask*. Although visual masking may call to mind a mask that spatially-overlaps the location of the target (e.g., pattern masking), masks can also affect perception when they never occupy a common spatial location with the target. Indeed, such forms have been more extensively studied in their own right, since the nature of the target perceptual interference is less obvious (Breitmeyer & Ögmen, 2006). OSM is a form of non-spatially-overlapping masking. In OSM, a target (e.g., Landolt C) is presented briefly with four small dots typically arranged to occupy the corners of an imaginary square centred on the target. Baseline target perception is measured when the target and mask appear and then disappear simultaneously (see Figure 1). Masking is then measured as the detriment in perception that results from the delayed disappearance of just the four-dot mask relative to this baseline condition. In other words, the temporally-trailing four-dot stimulus induces masking (Di Lollo et al., 2000; Goodhew et al., 2013).

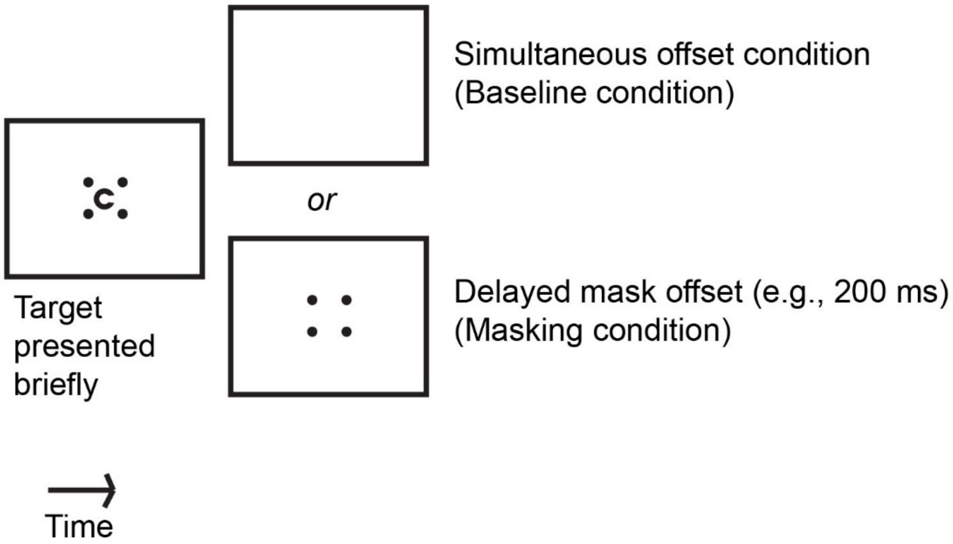

Figure 1. An illustration of object-substitution masking. The participants' task is typically to either detect or identify the target (e.g., in this case, to identify whether the gap in the C is on the left or right of the object). The target is presented briefly (<100 ms). Target perception is typically good in the simultaneous offset condition (where the target and mask disappear at the same time), whereas target perception is impaired when the four-dot mask stays visible for a fraction of a second after the target has disappeared relative to the simultaneous offset condition. This impairment is called object-substitution masking. These conditions can also be named according to the duration of the trailing mask (i.e., 0 ms for the simultaneous offset condition, and 200 ms for the delayed mask offset condition).

Of course, there is nothing special about the arrangement of the four-dot mask into a square, this simply reflects conventional use of the paradigm. Indeed, OSM is obtained when the target and four-dots are not aligned, or when the four-dots appear inside the contours of the target

(Guest, Gellatly, & Pilling, 2011), when the mask consists of only two dots presented on one side of the object (Kahan & Enns, 2010), and even when the mask is just a single dot (Lleras & Moore, 2003). In fact, even distractors (other non-target objects that are not the mask and typically look similar to the target) can drive masking, if their disappearance is delayed relative to the target (Luiga, Gellatly, & Bachmann, 2010). The key feature that characterises OSM is the continued presence of a non-target aspect of the array after the target has disappeared.

In the very first demonstration and report of OSM (Enns & Di Lollo, 1997), the target and mask were not invariably presented at the same time, and instead the interval of time between the target (stimulus onset asynchrony, SOA) was varied. While this set-up does indeed lead to masking, it is debatable whether this purely reflects OSM, or instead or a confounding mixture of OSM and metacontrast-masking (Breitmeyer & Ögmen, 2000). Therefore, virtually all subsequent studies in the literature have adopted the common-onset, temporally-trailing mask procedure to induce OSM. At the time, Enns and Di Lollo (1997) argued that the critical role of attention distinguished OSM from metacontrast masking, although as I will discuss later, it is now recognised that this not the case. For now, suffice to say that traditionally OSM studies have employed distractors (target-like objects that are presented in the target array in non-target locations) to manipulate spatial attention, but there is now a consensus that OSM can indeed occur for a single, centrally-presented and attended target (Filmer, Mattingley, & Dux, 2015).

Theories proposed to explain OSM: object substitution via re-entrant processing

The object-substitution model draws heavily upon the notion of re-entrant processing in the brain (Di Lollo et al., 2000). That is, rather than conceptualising perception as resulting from a single feedforward sweep progressing from posterior to anterior brain regions, models of re-entrant processing acknowledge the extensive feedback and iterative connections between different regions and their role in perception (e.g., Lamme & Roelfsema, 2000). In essence, the object-substitution model proposes that masking occurs when conflict between impoverished representations generated at higher levels in response to the briefly-presented target, and more high-resolution input reflecting the mask-alone stimulus is resolved by discarding the target representation in favour of the mask representation. That is, according to Di Lollo et al. (2000), the information from the initial array (target + mask) is encoded with low fidelity in the system, and then fed-forward from lower regions (e.g., V1) to higher regions (e.g., frontal). The low-fidelity of the representation is the consequence of at least two factors. First, the size of receptive fields of cells increases at higher levels, resulting in a spatially-coarse representation of the target at the higher levels, which will likely be insufficient for tasks such as target gap localisation, and second, it is possible for this initial representation to activate multiple potential representations at the higher level. This ambiguity and low-resolution representation need to be resolved, which triggers re-entrant processing returning to the site of the high-resolution content back in V1. Then a correlation process occurs, in which the putative representation generated is compared against the incoming sensory input in V1. If there is no change in the image, then a match is found, and processing continues. If the screen is blank, as it is in the simultaneous offset condition in OSM, then the system can use the decaying image and there is no competition from other, newly-incoming stimuli. On delayed mask offset trials, however, the stimulus has now changed – the target has disappeared and just the four-dots alone persist. This generates a new perceptual cycle and new tentative representation, and processing continues as per a new stimulus. Masking occurs when the perceptual representation containing the target is discarded in favour of the one containing only the mask (Di Lollo, 2010; Di Lollo et al., 2000). In addition to this description of the object-substitution process, Di Lollo et al. (2000) also provided a computational model to instantiate this process, called CMOS (computation model of object substitution).

The pivotal aspect of the object-substitution model to keep in mind is that it espouses a process of *replacement*, according to which there are separate object representations created for the target and mask, and the target representation is *discarded* during masking. In the original instantiation of this object-substitution account of OSM, it was hypothesised that either dispersing attention across the whole field or distracting focussed attention away from the location of the target is a necessary condition for substitution and therefore OSM. It was proposed that preventing focussed attention to the location of the target helped to create ambiguity about the identity of the target, thus fuelling the need for re-entrant processing to resolve the representation of the target (Di Lollo et al., 2000). I will return to discuss the role of attention in OSM more fully in later sections, and how this challenges the object-substitution account in a later section. But first, it is important to understand the object-updating account and how it is similar and how it is different from the object-substitution account.

Theories proposed to explain OSM: object-updating within the object-file framework

The object-updating account of OSM makes reference to the cognitive construct of object files. That is, the process of keeping track of individual objects in dynamic and busy scenes is challenging. Objects can move locations, the same object can look different due to changes in viewpoint or surrounding objects, and we have frequent interruptions to visual processing during saccades. Yet despite this, the process of tracking object identities across time is one that occurs largely seamlessly for those with healthy, functioning visual systems. This apparent seamlessness is attributed to mid-level representations called *object files* or *object tokens* (Kahneman & Treisman, 1984; Kahneman, Treisman, & Gibbs, 1992; Kanwisher, 1987, 1991; Kanwisher & Driver, 1992). Kahneman et al. (1992, p. 178) describe an object file as follows:

> "Each object file is addressed by its location at a particular time, not by any feature or identifying label. It collects the sensory information that has so far been received about the object at that location. This information can be matched to stored descriptions to identify or classify the object, but it need not be. We can normally see completely novel objects with little difficulty, without knowing what they are. When the sensory situation changes, the information in the files is updated, yielding the perceptual experience of changing or moving objects. A file is kept open so long as its object is in view, and may be discarded shortly thereafter. The system bridges over the discontinuities produced by temporary occlusion, or by saccades, assigning current information to preexisting files whenever possible"

A key point from this definition is that the identity of the object in a given file is not fixed, but can be updated over time. Kahneman et al.'s (1992) classic work compellingly demonstrated that two static presentations of arrays separated in time could be linked by such object-files. That is, when the stimuli reappeared in new locations that were consistent with them as objects having moved along a consistent spatiotemporal trajectory, participants were more adept at perceiving and responding to the stimuli than when their positions violated a consistent spatiotemporal trajectory (Kahneman et al., 1992). More recent research suggests that consistency in basic surface features (e.g., colour) over time also contribute to the maintenance of object files, in addition to consistency in spatiotemporal trajectory (Goodhew, Fogel, & Pratt, 2014; Hein & Moore, 2012; Hollingworth & Franconeri, 2009; Richard, Luck, & Hollingworth, 2008). For OSM, this means that if the visual systems treats the trailing four-dot mask as a continuation of the previous array (which contained the four-dots and also the target), then we would expect the object representation containing the target to be updated to reflect the four-dot mask alone. Indeed, not only is this possible, but according to the object-file perspective, the system prefers it ("the system bridges over discontinuities…assigning current information to pre-existing files *whenever possible*" [emphasis added]). This is highly pertinent to OSM: if the system favours the modifying object representations to reflect the current visual scene, then it is more likely that there is a single object representation, which initially reflects the target plus the four-dots, which, under conditions of delayed mask offset, is subsequently updated to reflect just the four-dots alone.

This is the crux of Lleras and Moore's (2003) *object-updating* account of OSM: there are not separate object representations of the target and mask, but instead a single representation which is *updated* over time. This updating means that the initial representation that contained information about the target is *modified* or over-written, thus resulting in impaired target perception. Specifically, in the delayed-mask offset condition, the representation is modified to reflect the four-dot mask alone and this is the end-state of the object-file, and thus only the four-dot mask is consciously perceived. Whereas in the simultaneous mask offset condition, the object representation that initially represents the target is not updated beyond this because there is no conflicting input, and thus the end-state of the object file is the representation containing the target, and therefore the target is consciously perceived.

Object Substitution vs. Object Updating

At first blush, the distinction between the object-substitution and object-updating accounts may appear to be hair-splitting trivialities at best, or unfalsifiable at worst. After all, the functional outcome of the two processes on masking is identical: mask consciously perceived, target lost from conscious perception. Evidence of unconscious target processing (e.g., Binsted, Brownwell, Vorontsova, Heath, & Saucier, 2007; Goodhew, Visser, Lipp, & Dux, 2011), also does not distinguish the accounts, as both object-substitution and object-updating allow for initial representation, and therefore processing, of the target. But the distinction between the two accounts is actually fundamentally important, due to the broader frameworks to which these accounts belong. The object-file framework and object-updating account make clear and direct predictions that the object-substitution account does not. That is, understanding OSM as part of the broader visual-cognitive processes of assigning and tracking object identities over time and change makes specific predictions that do not arise from object-substitution via re-entrant processing framework, and these predictions as we shall see, are supported by a plethora of evidence.

To reiterate, according to the object-updating account of OSM, the fact that the target is not consciously perceived results from a (non-conscious) perceptual inference that the trailing-mask (second array) reflects a continuation of the same object that began with the target plus four-dot mask (first array). This means that masking is not the result of any fundamental temporal processing limitation or bottleneck, but instead the result of an adaptive inference that allows for the continuation of object identities over time. From this it follows that the masking should be flexibly influenced by factors that influence inferences of object continuity. More specifically, if the object-updating account is true, then, manipulations that encourage the visual system to treat the target and four-dots as part of the same object should increase the likelihood of object-updating and therefore increase masking magnitude. Manipulations that encourage the visual system to infer that the target and four-dots to be individuated as distinct objects, should discourage object-updating and therefore reduce masking. In contrast, the object-substitution account makes no such predictions about inferences of object-continuity versus object-individuation.

There are many demonstrations of factors that are known to influence the likelihood with which the system will treat the target and trailing mask as the same object systematically influence the magnitude of OSM. In the seminal demonstration of this, Lleras and Moore (2003) showed that what happens to the display *after* the mask has initially disappeared can influence whether masking was obtained or not. Specifically, they had four-dot masks appear around all of the items in an annulus target array, both the target and distractors. The target was distinguished from the distractors by appearing in a different colour. This set of four-dot masks disappeared simultaneously with the target, but then reappeared at an outer location after either a short interval or a long interval. When the interval was short, it was conducive to apparent motion, such that it appeared as though the annulus of four-dot masks from the target array subsequently expanded into a larger-circumference annulus. When the interval was long, however, it did not facilitate the perception of apparent motion, and thus would appear as though a new annulus appeared (see Figure 2).

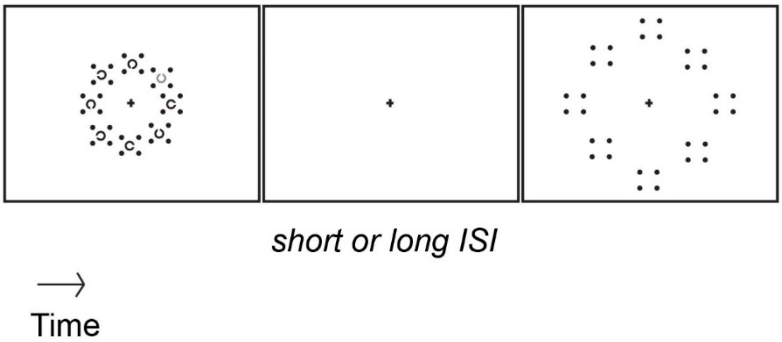

Figure 2. This shows the critical frames from Lleras and Moore (2003). In the target display on the left, the four-dot masks for each of the objects are arranged in a circle, which expands in circumference during the trailing-mask array on the right. The critical manipulation, therefore, is whether the interstimulus interval (ISI) between the target and trailing-mask array is short, such that the latter appears to be an instance of the former in a new location (linked via apparent motion), or long, such that the trailing-mask array appears to be the onset of a new object, unrelated to the target array.

Target-identification performance in these two conditions was compared against a baseline condition in which the four-dot masks disappeared simultaneously with the target array and did not reappear. Robust masking was found in the short-interval condition, and was eliminated in the long-interval condition. That is, when the interval between the disappearance of the target array and the appearance of the mask-only array encouraged the inference that the larger-circumference annulus was a continuation of the stimuli from the target array, masking occurred, whereas when the duration of this interval discouraged this object-continuity inference and instead encouraged the inference of a new object appearing, masking did not occur. This is consistent with the object-updating model of OSM, according to which when the trailing mask is perceived as a modified continuation of the target array, the object representation that includes the target is updated to reflect the state of the trailing-mask alone and this is how information about the identity of the target is lost. The object-substitution account, in contrast, struggles to explain this finding. Since there was no immediately temporally-trailing four-dot mask, there should not have been new input to interfere with the perceptual iterations required to render the target representation consciously perceived.

The only way that the object-substitution account could be salvaged in light of Lleras and Moore's (2003) finding, is if in the short-interval condition there was greater interference in the iterative loops required to resolve the target identity than in the long-interval condition. This seems a problematic interpretation, however, given that the trailing mask are in different spatial locations, and thus unlikely to be the focus of re-entrant processing to refine the target representation. However, subsequent follow-up studies rule out even this post-hoc interpretation. It has been shown that the presentation of the annulus of dots in an outermost location in between the two presentations in the short-interval condition – which should interfere with the percept of the annulus mask array as a continuing object travelling along a consistent trajectory – significantly reduced masking the short-interval condition (Pilling & Gellatly, 2010). Such a presentation should have either no bearing on the process of object-substitution (due to the large spatial separation), or a tendency to increase substitution and therefore masking by increasing exposure to the mask-array excluding the target in the critical time window after the brief presentation of the target. But masking was reduced. Thus object-substitution theory can offer no explanation of this effect, which the object-updating account can explain.

Furthermore, since this initial demonstration in support of the object-updating account, there have been subsequent findings which were inspired by this framework, and for which the object-substitution account cannot offer any plausible interpretation. Another body of findings that favour object-updating over the object-substitution account is that static featural factors that make the target and four-dot mask appear part of the same object increase masking, whereas these same factors that make the target and four-dot mask appear part of different objects decrease masking. Namely, when the target and four-dot mask appear in the same colour, luminance polarity, or orientation as one another, masking is increased, whereas when they differ along these dimensions, masking is reduced (Goodhew et al., 2015; Luiga & Bachmann, 2008; Moore & Lleras, 2005). Such manipulations are ideally suited to encouraging and discouraging the inference that the target and trailing mask belong the same versus independent object representations, and therefore the object representation containing the target should be updated versus a new one created to reflect the trailing mask. According to object-substitution, however, such manipulations should have no discernible effect. This is because the correlation between the descending signal containing the target plus four-dot mask and the incoming input representing the trailing mask alone should be constant regardless of whether the target and four-dot mask are similar versus different, it would only be influenced by whether the trailing four-dot mask and preceding four-dot mask in the target-array were similar versus different.

Research has also revealed that *preview* of the four-dot mask prior to the presentation of the target array reduces OSM (Neill, Hutchison, & Graves, 2002). This is true even when the target appears amongst distractors and the location of the mask preview does not predict the location of the target, and thus the advantage in target perception cannot be a consequence of attentional cueing. The object-updating framework explains this phenomenon as reflecting the creation and consolidation of a unique object-representation for the four-dot mask, thereby reducing the likelihood that the target representation will be updated to reflect the temporally-trailing mask. Furthermore, in follow-up examinations of the mask-preview effect, it has been shown that a change in the configuration or shape of the mask (e.g., from square of four-dots to diamond of four-dots), the preview effect was diminished, except when the interval of time separating the preview from the target array mask was 0 ms, thereby presumably giving rise to the percept of apparent motion between the two and thus they would be treated as the same object (Lim & Chua, 2008). One limitation of the Lim and Chua (2008) study, however, is that they only tested the effect of their manipulations on raw accuracy, rather than masking magnitude (difference in accuracy between simultaneous and delayed mask offset conditions). This means that we cannot be sure that these manipulations truly altered masking, rather than generic target perception per se.

Subsequent research has shown, however, that there is a selective reduction in OSM magnitude due to the benefit of object preview. Specifically, previewing the targets and distractors prior to the onset of the four-dot mask reduced masking (Gellatly, Pilling, Carter, & Guest, 2010), even when the items that are previewed are simple placeholders (e.g., squares) which do not reveal anything about the stimulus identity required for report (e.g., gap on left versus right of object) (Guest, Gellatly, & Pilling, 2012). This is strong support in favour of the object-updating explanation of OSM. Object-substitution, in contrast, has no provision for stimulus displays prior to the onset of the target affecting masking, unless they serve to cue the location of the target, an advantage which neither the preview of target array stimuli in the absence of the mask (Gellatly et al., 2010), nor mere placeholders (Guest et al., 2012), offer. In other words, the object-substitution account is notably silent on how such preview effects can moderate OSM. These effects are fully consistent, in contrast, with the object-updating framework, because they affect the likelihood of the system inferring whether the four-dot mask reflects a continuation of the target array, and therefore whether the target representation ought to be updated.

Moreover, another pivotal difference between the putative mechanisms of object-substitution and object-updating, is whether they permit any perceptual interactions between the target and mask representations. The object-substitution framework stipulates separate representations for the target and mask – and therefore while the mask can affect whether or not the target representation is discarded, this is a dichotomous all-or-nothing process: it is either discarded or not, and no intermediate states or incomplete mergers are possible. Under the object-updating framework, in contrast, there is a single representation which contains both the target and mask at different points, and is modified from the target to the mask in the face of the trailing-mask array. Logically, this updating process could be incomplete, and since information about the target and mask have been contained within the same file, it raises the possibility of perceptual interactions emerging. Indeed, such interactions have been observed. OSM can alter the *nature* (Kahan & Enns, 2010) and the *quality* of the representation of the target (Harrison, Rajsic, & Wilson, 2016), rather than entirely eliminate the target representation. These findings refute the theory of object-substitution.

The role of attention in object-substitution masking

What is the role of attention in OSM, and what does it tell us about the underlying mechanism? In the early days, the interaction between set-size (the number of items in the target array) and trailing mask duration was touted as a hallmark of OSM, such that masking is stronger when attention is distributed over more locations. Increasing the set-size of the array was assumed to distribute visual attention, preventing focussed attention on the target. Furthermore, OSM was believed to be absent when only the target was presented alone without the presence of distractors (Di Lollo et al., 2000). Indeed, I have made statements to this effect (Goodhew, Dux, Lipp, & Visser, 2012). But, reflecting on more recent evidence that has come to light, I think that such statements place too much emphasis on the interaction between attention and OSM. To be clear, set-size can and does have a dramatic effect on overall target perception. The issue in question, however, is whether this increases the *magnitude of the masking* (i.e., the difference in accuracy between simultaneous and delayed mask offset conditions) in particular. Argyropoulos, Gellatly, Pilling, and Carter (2013) were the first to call into question the role of attention in OSM, arguing that putative interactions between set-size and mask duration reported previously were the spurious consequence of floor and/or ceiling effects, and demonstrated that when these factors were corrected, the interaction was absent (see also Filmer, Mattingley, & Dux, 2014). Furthermore, it has also been convincingly demonstrated that OSM can occur for a single, centrally-presented target, for which there is no place for the distribution of attention (Filmer et al., 2015). However, before we completely dismiss the role of attention as having no business in OSM studies, Camp, Pilling, Argyropoulos, and Gellatly (2015) noted that in the previous studies that obtained the absence of an interaction between set-size and masking, there was considerable spatial distance between the items in the target array. When these authors reduced the distance down into the critical zone conducive to visual crowding (Bouma, 1970; Greenwood, Bex, & Dakin, 2010; Pelli & Tillman, 2008), the interaction was now reliably obtained (Camp et al., 2015).

This highlights a potential problem with manipulating *attention* via the presence or absence, or the number of items in the array: this potentially introduces a number of other artefacts that are not specifically related to attention. For example, the potential for crowding (Pelli & Tillman, 2008), and/or perceptual dilution (Tsal & Benoni, 2010) is present. For this reason, it is important to have a manipulation of attention that does not introduce these issues to provide converging evidence regarding the role of attention in OSM. Changing the *size* of the attended region affords this opportunity. That is, humans can regulate the size of their attended region, in the absence of a shift of the central location of attention (LaBerge, 1983), such that when resources are applied over a small region there is a greater magnitude of perceptual enhancement within this attended region than when resources are spread over a larger region (Eriksen & St. James, 1986; Goodhew, Shen, & Edwards, 2016; Muller, Bartelt, Donner, Villringer, & Brandt, 2003). Recent research demonstrates that while such changes reliably impact spatial acuity, they have *no* impact on the magnitude of OSM (Goodhew & Edwards, 2016). This is converging evidence that OSM does not interact with visual-attentional processes (see Figure 3).

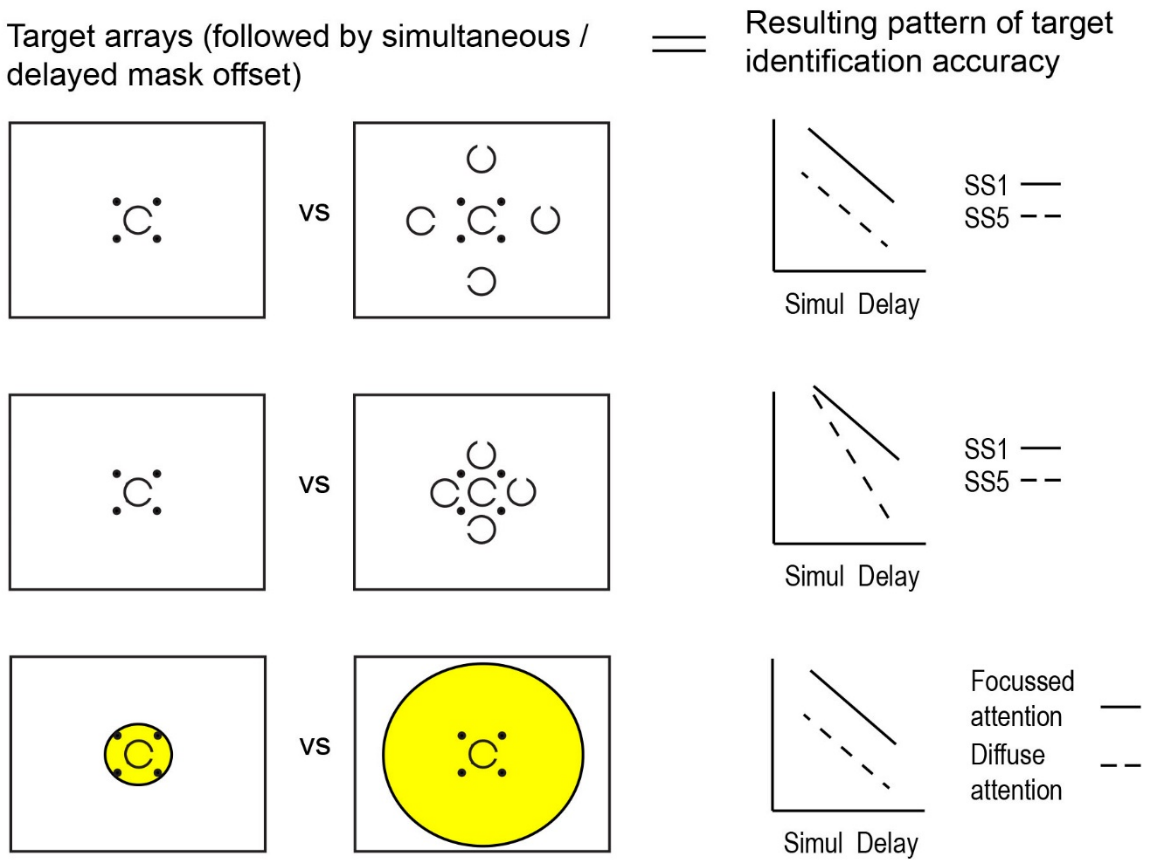

Figure 3. An illustration of the target arrays and the pattern of results they generate when followed by simultaneous versus delayed mask offset frames (see Figure 1). The y-axis for the indicative results graphs is target identification accuracy. In the top row, there are both masking

and set-size effects, but these do not interact when there is wide target-distractor spacing (Argyropoulos et al., 2013), whereas in the second row, they do interact with narrow target-distractor spacing, such that masking is exacerbated at larger set-sizes (Camp et al., 2015). Finally, in the third row, a narrow versus a wide attended-region impacts spatial acuity, but does not interact with masking magnitude (Goodhew & Edwards, 2016). The yellow area indicates presumed span of spatial attention in response to manipulation. The fact that the interaction is not present for an attentional manipulation but is present for tightly packed arrays suggests that the interaction with the latter is a product of crowding. Simul = simultaneous target and mask offset, delay = delayed mask offset. SS1 = set-size 1 (i.e., target-only array), SS5 = set-size 5 (i.e., target plus four distractors).

To summarise, the evidence indicates that the distribution of attention is not a necessary requirement for OSM. Neither is the interaction between set-size and mask duration a defining hallmark of OSM. However, it is also not the case that the interaction *never* occurs in OSM – it can, but just under certain conditions, such as the spacing between the items creating visual crowding. However, this appears to the case when mechanisms *other* than attention are available to explain the interaction (e.g., crowding), and that when these are removed, so is the interaction.

The original description of the object-substitution model emphasised that the distribution of attention was a key aspect in creating ambiguity about the perceptual hypothesis generated in frontal areas and therefore susceptibility to masking (Di Lollo et al., 2000). Inferences of object-updating versus object-individuation, in contrast, have not been hinged on preventing visual attention applied to the stimulus in question. In this sense, the recent evidence to come to light that attentional manipulations are not at the heart of OSM could be considered further evidence to substantiate the object-updating over the object-substitution account.

That said, a counterargument to this is that while set-size may have been a feature of the original description of OSM, the essential machinations of substitution can be salvaged by updating the account to omit the role of attention. Indeed, rapid presentation time and the increase in receptive field size in the higher regions responsible for generating the perceptual hypothesis are sufficient in their own right to explain the need for re-entrant processing in the face of an OSM display. Put another way, object-substitution was likely linked to attention in the first place to explain the presence of set-size effects observed in the data, but the fundamentals of object-substitution do not depend on it. Together, the findings that dynamic and static featural manipulations that affect inferences about object identity also affect OSM, the fact that preview of placeholder stimuli can mitigate masking, and the clear evidence that the target representation can be altered without be discarded all much more strongly controvert the object-substitution account than the newly-understood role of attention. That said, the current state of play in the literature is clearly that OSM and attention are best conceptualised as independent processes.

Can substitution and updating mechanisms together be implicated in OSM?

So far I have emphasised pitting the substitution account versus the object-updating account in and ‘either/or’ fashion. An alternative way of thinking about it, however, is that both processes could be implicated, either in concert, or independently at different points in time. This is a possibility, and one that no amount of evidence in favour of object-updating can refute. At the very least, however, the purpose of this commentary is to support the conclusion that OSM unambiguously implicates object-updating and can be used to investigate such processes. Moreover, while logically it cannot be dismissed, I do not favour the conclusion that substitution processes are implicated in OSM in a meaningful way, because despite a burgeoning OSM literature, convincing evidence for substitution processes that cannot be explained by updating processes remains to be seen. Notably, Kahan and Lichtmann (2006) explicitly sought to test whether substitution processes were involved in OSM, in addition to the updating processes previously documented in Lleras and Moore (2003). Kahan and Lichtman (2006) presented the target and mask in various depth planes, and found that masking occurred when the target and mask were presented in the same depth plane, regardless of whether near or far, but when they were presented in different depth planes, masking only occurred when the mask appeared in front of the target, not vice versa (Kahan & Lichtman, 2006). Kahan and Lichtman (2006) suggested that this is evidence for separate object tokens in competition, rather than a single file that is updated. However, there are two issues with this interpretation. First, it is difficult to clearly interpret what these results mean, since baseline performance for the target-in-front condition was worse than performance for the other conditions, and masking magnitude (by virtue of being a difference score) can be constrained by baseline accuracy. Second, even if we take the masking attenuation in the target-in-front condition at face value, then this result can also be explained within an updating framework, since updating will simply be more likely to occur for a less consolidated target (as evident from the demonstrations that placeholders protect against masking). If objects in the foreground are processed first relative to those in the background, then the target will simply be less consolidated and therefore more vulnerable to masking when the mask is in front, compared with when the target is in front. Furthermore, Kahan and Lichtman's (2006) subsequent demonstration that a mask moving toward the target produced more masking than one moving away does not appear to cleanly adjudicate between the two accounts, since the movement toward could encourage updating to a greater extent than movement away. Altogether, while conceptually it remains possible for substitution and updating processes to both play a role in explaining OSM, the weight of the evidence is clearly in favour of object-updating.

The relationship between OSM and other visual-cognitive phenomena

The visual-cognitive phenomena with which OSM does and does not interact speaks volumes about how its driving mechanism should be conceptualised. When two phenomena interact, it points to shared mechanisms shared between the two, whereas if two phenomena do not interact but remain separate, then it is evidence to suggest that they do not have common underlying mechanisms. In this vein, OSM interacts with a phenomenon that sit firmly in the object-individuation processes camp (akin object-updating), and fails to interact with a phenomenon that is thought to reflect temporal processing bottleneck (akin to object-substitution). I will also discuss other phenomena in where the relationship with OSM has not been directly tested, but the conceptual relationship between the two is theoretically important. I will also highlight where OSM fits in ongoing debates with related phenomena, such as object correspondence through occlusion.

Repetition blindness. OSM interacts with repetition blindness (Goodhew, Greenwood, & Edwards, 2016). Repetition blindness refers to the discovery that repeating the identity of a stimulus close in time can impair the individuation of a unique episodic representation for both instances (Kanwisher, 1987). This means, for example, that if two instances of the letter 'A' are presented either simultaneously or in quick succession, then participants will typically recognise that a letter A has appeared ('type recognition'), but form just one instead of two distinct episodic representations ('tokens') of that letter A. That is, people are 'blind' to the repetition (Bavelier, 1994; Kanwisher, 1987, 1991; Kanwisher, Kim, & Wickens, 1996; Kanwisher & Potter, 1989). Repetition blindness is strongly modulated by factors that manipulate the episodic distinctiveness of the items despite identical timing and stimulus duration parameters (Dux & Coltheart, 2008; Goldfarb & Treisman, 2011). Repetition blindness therefore fundamentally gauges how the visual system assigns object identities over time in the face of rapid input, that is, inferences of object-updating versus object-individuation. The fact that repetition blindness and OSM interact (Goodhew et al., 2016) is therefore further evidence in favour of the object-updating account of OSM.

Attentional blink. In contrast, OSM does not interact with the attentional blink (Giesbrecht, Bischof, & Kingstone, 2003). The attentional blink refers to the detriment in perception of a second target in a rapid, serial, visual presentation stream that occurs approximately 200-500 ms after identifying the first target (Raymond, Shapiro, & Arnell, 1992). The general consensus is that the attentional blink reflects an intrinsic temporal bottleneck in processing and consolidating items into visual short-term memory (Chun & Potter, 1995; Dell'Acqua et al., 2015; Dell'Acqua, Dux, Wyble, & Jolicoeur, 2012; Dux & Marois, 2009; Vogel, Luck, & Shapiro, 1998). This is akin to object-substitution theory, which attributes to masking to an intrinsic inability to extract and consolidate to the level of awareness given the spatiotemporal properties of the OSM display. The fact that OSM does not reliably interact with the attentional blink (Giesbrecht et al., 2003), therefore, is evidence against the object-substitution account of OSM.

Flash-lag effect. The flash-lag effect refers to an illusion in which a moving object appears to be spatially “ahead” of a static object, despite the fact that they are actually veridically aligned (Eagleman & Sejnowski, 2000; Nijhawan, 2002). More specifically, in arrays designed to elicit this illusion, a disc moves along a continuous circular trajectory, and a small static square briefly flashes up. If the disc continues moving, then participants report perceiving that at the moment that they appeared simultaneously, the moving disc appeared spatially ahead (further along its circular trajectory) compared with the static square. In contrast, if the disc’s motion ceases at the time that the static square appears, then the illusion is eliminated, and instead observers perceive the disc and square as having been aligned.

Enns, Lleras, and Moore (2010) marshalled convincing evidence that object-updating occurs in the service of creating perceptual continuity and stability of visual scenes. That is, the visual system draws a sharp distinction between when an object is considered to continue in a scene, in which case information about the object (e.g., location) is updated, versus when a new object appears in the scene and this updating does not occur. Enns et al. (2010) describe a variety of phenomena and demonstrations with them that speak to the importance of object-updating, including a particular focus on the flash-lag effect. That is, research has shown compellingly that this flash-lag illusion can be explained within an object-updating framework, whereby information in dynamic scenes is integrated into existing representations of the scene (Enns et al., 2010; Moore & Enns, 2004). This means that if information is seen as coming from a continuing object, it gets updated to reflect more recent information at the expense of the earlier information, whereas it is protected from this process if not treated as part of the same object. This was shown convincingly when it was demonstrated that abrupt changes in the features of the travelling object (size, colour) at the time of the flash eliminated the flash-lag illusion even when the disc continued to travel after the flash (Moore & Enns, 2004). This supports the object-updating account of the flash-lag effect, because such manipulations would discourage treating the disc as a continuous object pre- and post-flash. Most strikingly, in these instances where an abrupt change eliminated the flash-lag illusion, participants were more likely to report seeing *two* separate discs, when in fact there was only one. This is because the transient change induced a new object file, and therefore encouraged an illusory percept of an additional disc. This is strong support for the notion that the flash-lag effect is a product of object-updating processes (Enns et al., 2010; Moore & Enns, 2004).

Conceptually, therefore, OSM and the flash-lag effect essentially reflect the same object-updating mechanism, which allows for object continuity to occur in dynamic scenes. The continued motion of the disc in the flash-lag effect is akin to the delayed offset of the four-dot mask in OSM. Both of these stimuli alter what is consciously perceived, and both illusions are broken by featural changes and spatiotemporal interruptions. It would be interesting to test the similarity between the two mechanisms further, for example, by examining whether individuals who are most susceptible to OSM, are also more susceptible to the flash-lag effect, but not any more or less susceptible to an unrelated phenomenon, such as the AB.

Object continuity across interruption in visual input. There are many instances in which visual input from an object is interrupted. This occurs during saccades, a process during which humans suppress visual input and are functionally blind (Ibbotson & Krekelberg, 2011; Matin, 1974), or by having an object temporarily disappear behind another object and thus is physically occluded (Burke, 1952; Flombaum & Scholl, 2006). Furthermore, objects can move (i.e., change location), and their features can also be altered, due to changes intrinsic to the object itself, or due to viewpoint variation (e.g., seeing the same person from different angles). An important question, then, is how does the visual system maintain continuous object representations despite such changes? What properties help the brain to tell when an object continues versus when a new one appears? Historically, such questions were addressed with the object-specific preview paradigm (Kahneman et al., 1992), from which the theoretical construct of object files were developed. The work showed that presentations of physically-identical stimuli in different spatially locations separated by a brief interval of time were linked together as an object continuing over time based on their *spatiotemporal history*. In other words, if they appeared at spatial locations and in the timeframe consistent with the two presentations showing the objects initially in one location, and then having moved to another.

More recently, such questions have been addressed by explicitly showing objects moving behind other objects (object correspondence through occlusion), or by examining how participants respond to changes in location or features (e.g., colour) of objects during saccadic suppression. For example, if participants initiated a saccade toward an object, and then it moved location, or changed colour – would they still perceive it as the same object? (e.g., Richard et al., 2008). Initially, the primacy of an object's spatiotemporal trajectory was highlighted (Flombaum, Kundey, Santos, & Scholl, 2004; Flombaum & Scholl, 2006; Yi et al., 2008), and this led to the conclusion that spatiotemporal trajectory is the sole or critical determinant of whether a continuous versus a discontinuous object is perceived over time. This means that it would be of no consequence if an object change features such as shape, colour, or orientation in front of you – provided that travelled along a consistent path through space and time, you would see it as the same object. Subsequent research has refuted this conclusion, and convincingly shown that an object's surface features also contribute to stable object representations, in addition to spatiotemporal trajectory (Hein & Moore, 2012; Hollingworth & Franconeri, 2009; Moore, Stephens, & Hein, 2010). For example, changing either the position or the colour of an object during saccadic suppression interferes with the percept of a continuing object over time (Richard et al., 2008). This is important, because this conclusion converges with that arising from OSM, that both spatiotemporal trajectory and surface features affect whether or not object-updating and therefore masking occurs (e.g., Lleras & Moore, 2003; Luiga & Bachmann, 2008; Moore & Lleras, 2005).

Why, then, did the initial studies appear to support the contradictory conclusion of spatiotemporal dominance in object-updating processes? In fact, closer examination of the evidence that was purported to support the spatiotemporal dominance hypothesis actually reveals that both spatiotemporal and surface features co-contributed all along. For example, Flombaum, Kundey, Santos and Scholl (2004) tested the role of surface features and spatiotemporal trajectory to object correspondence in macaque monkeys (Macaca mulatta). To do this, they recorded the monkeys' search behaviour after watching a moving fruit object (e.g., a lemon) roll down a ramp and disappear behind one occluder, and a new object emerge (e.g., a kiwifruit), and subsequently roll down behind a second occluder. The monkeys were exposed to the following four conditions: tunnel events (in which one fruit disappeared behind an occluder, and another, different fruit emerged, i.e., featural change, but on a continuous spatiotemporal trajectory), temporal gap with featural change (the fruits changed, and had a three-second pause behind the occluder before it emerged), temporal gap without featural change (paused for three seconds, but then the same fruit emerged), and simultaneous presentation, in which two different fruits were exposed simultaneously. The authors relied on the frequency of monkey's subsequent search behind the first occluder as a measure of object discontinuity. This is because if the monkeys perceived a new object as emerging from behind the occluder, they should search behind the first occluder to locate the fruit obscured there. So the perception of a continuing object over time is inferred from decreased search behind the first occluder (Flombaum et al., 2004).

Flombaum et al. (2004) found decreased search behind the first occluder in the tunnel condition (21%), compared to temporal-gap-with-feature-change (61%, or the simultaneous presentation condition, 65%), indicating that the temporal gap disrupted continuing object identity (Flombaum et al., 2004). However, in comparing the two conditions that both had a temporal gap, search behind the first occluder was increased when there *was* a featural change (61%) compared to when there was *not* and the objects were featurally identical (15%), indicating that a change in ***features*** also disrupted continuing object identity while spatiotemporal trajectory was held constant. Furthermore, there were equivalent search rates for tunnel-event condition (change in features but uninterrupted spatiotemporal trajectory), as for the temporal-gap-without-featural change (constant features but interrupted temporal trajectory). This means that a change in *either* property – features or trajectory – had a similar effect on object continuity perception. While the authors emphasised the role of spatiotemporal continuity (e.g., "In this study we demonstrated the tunnel effect – and its associated spatiotemporal bias – for the first time in a nonhuman primate" (Flombaum et al., 2004 p. 799), objectively the results clearly speak to at least an equal role of featural continuity in determining object persistence through occlusion. Similarly, other evidence that has been proposed to support the conclusion that spatiotemporal trajectory trumps features in humans in determining object correspondence have typically only examined the influence of spatiotemporal changes on objects' feature changes (Flombaum & Scholl, 2006), or on object-level cortical responses (Yi et al., 2008) – thereby not permitting any conclusions about *relative* contribution of spatiotemporal changes versus features to be drawn.

How does this relate to OSM? If the spatiotemporal dominance hypothesis were true, then this would put OSM at odds with the other literature on how object representations are created and maintained across time. This is because OSM has shown that masking magnitude (and by inference, the formation of single versus multiple object representations) is influenced by both spatiotemporal trajectory (e.g., how the mask moves after target offset), and features (e.g., whether the target and mask appear in same versus different colours). This means that as a consequence of re-examining and rejecting the spatiotemporal dominance hypothesis, there is now converging evidence from multiple paradigms that the formation of object files is influenced both by how an object looks, as well as where it has been and where it is going, and when.

Feature inheritance, backward masking, and apparent motion. Feature attribution is where the second of two stimuli subjectively 'inherits' the objective features of the first. In contrast, in both apparent motion and backward masking (where the mask appears after the target, unlike the common onset in OSM), the second of two stimuli impairs perception of the first stimulus. Breitmeyer, Herzog, and Ogmen (2008) examined the relationship between feature attribution, backward masking and apparent motion. The fact that feature attribution and backward masking occur under similar conditions has been attributed to shared theoretical mechanisms (Breitmeyer, Herzog, & Ogmen, 2008; Enns, 2002; Herzog & Koch, 2001). But empirically, Breitmeyer et al. (2008) found that apparent motion and feature attribution had the strongest relationship, stronger than that between backward masking and feature attribution. Moreover, the modest relationship between backward masking and feature attribution could be entirely explained by the relationship between apparent motion and feature attribution. This suggests that apparent motion and feature attribution belong to the broader body of object-updating mechanisms. Given the evidence that OSM also belongs to this category, it suggests that OSM and backward masking are distinct. It therefore remains possible that object-substitution is a viable explanation for backward masking, whereas recourse to object-updating is required in order to understand OSM.

Motion-induced blindness (MIB). MIB is the striking alteration in visual awareness that occurs whereby static targets that are presented continuously (e.g., for 30 seconds) subjectively 'disappear' from awareness when presented in close proximity to moving elements (Bonneh, Cooperman, & Sagi, 2001). Research has shown that MIB is enhanced at the trailing edge of movement relative to the leading edge (Wallis & Arnold, 2009). This suggests that MIB is not a maladaptive 'failure' of visual awareness, but instead likely reflects the adaptive mechanism that suppresses motion streaks (the processing 'smear' that trails behind a moving object) from our visual awareness (Wallis & Arnold, 2009).

What is the relationship between OSM and MIB? The perception of motion implies the presence of a continued object representation over time – an object whose location is being updated, rather than a new object representation being formed in response to stimulation at new locations. Motion mechanisms have already been established to play a role in OSM. For example, the nature of apparent motion of the trailing mask can influence whether masking is observed, dependent on whether the mask is perceived as continuing object or not (Lleras & Moore, 2003), and repetitive transcranial magnetic stimulation (rTMS) applied to motion-sensitive area V5/MT+ attenuates masking (Hirose et al., 2007). This suggests that the normal object-updating mechanisms that are implicated in motion are also implicated in the suppression of the target that typifies OSM. MIB and OSM have in common a selective suppression of visual processing that appears to have adaptive purpose. If there was not suppression of motion streaks our perception of the world would be blurry and streaky, and if we did not update object representations over time, every change or movement in an object would be perceived as an entirely new object and there would be no continuation of objects over time. But what is suppressed is defined in different ways. For MIB, it depends on the spatial arrangement of the array (what is leading/trailing in space), whereas for OSM it depends on the temporal arrangement of the array (what is leading/trailing in time).

Inattentional blindness (IB). IB refers to the failure of visual awareness (i.e., failure to detect presence) of an unexpected object due to attentional resources being otherwise occupied (Kim & Blake, 2005; Rock, Linnet, Grant, & Mack, 1992). IB increases with increasing perceptual load (Cartwright-Finch & Lavie, 2007), and one of the ways in which perceptual load is manipulated is number of search items (i.e., akin to set-size in OSM). In contrast, the evidence suggests that OSM is not affected by set-size (Argyropoulos et al., 2013), unless crowding (Camp et al., 2015), rather than the diffusion of attentional resources (Goodhew & Edwards, 2016), is implicated as the underlying mechanism. This points to different mechanisms behind OSM and IB.

Furthermore, while impaired target perception in OSM appears to be on a graded level – there still remains advanced processing despite effective masking (Goodhew et al., 2011), and masking can alter the *quality* of the representation rather than eliminate it altogether (Harrison et al., 2016; Kahan & Enns, 2010), in IB, the evidence to date suggests more of a rigid bifurcation in whether or not the target is consciously perceived. This suggests that IB might reflect more about central bottlenecks which are harmful limitations to processing that can be dangerous (Most, Scholl, Clifford, & Simons, 2005), rather than adaptive visual-cognitive mechanisms (Goodhew, 2016). Altogether this supports the notion that the IB reflects dissociable mechanisms from OSM.

Biological basis of object updating

Extensive behavioural evidence implicates inferences of object continuity in OSM. That is, the evidence supports the object-updating account, according to which the target representation is updated to reflect the trailing four-dot mask. Moreover, a number of other related visual-cognitive phenomena also reflect the operations that permit continuing object representations over time and change, including repetition blindness, apparent motion, feature attribution, and the flash-lag effect. Given this, it is worthwhile considering the putative neural basis to these cognitive and behavioural effects. The neuroscientific basis of object continuity as gauged through apparent motion has been well-researched. Evidence tells us that apparent motion activates area V5/MT+ sensitive to real physical motion (Mikami, 1991; Muckli et al., 2002), and motion implies the continuation of an object representation over time and change. Moreover, V1 is activated in the regions that retintotopically lie between the two locations activated by the two presentations of stimuli that induce apparent motion (Muckli, Kohler, Kriegeskorte, & Singer, 2005). This suggests that re-entrant processing originating from V5/MT+ perceptually “fills in” the activity in V1 that corresponds to the path of the apparent motion (Muckli et al., 2005). While it remains to be directly tested, it is possible that such a re-entrant mechanism could be the basis of other object-updating phenomena, including OSM. Indeed, this is broadly consistent with the findings that rTMS to V5/MT+ thwarts OSM (Hirose et al., 2007), and that re-entrant activation of V1 is implicated in OSM (Boehler, Schoenfeld, Heinze, & Hopf, 2008; Kotsoni, Csibra, Mareschal, & Johnson, 2007).

Conclusion

The object-substitution account of OSM stipulates that there are two separate object representations, one for the mask and one for the target, and that ultimately masking occurs when the mask representation replaces the target. The object-updating framework, in contrast, espouses that OSM reflects an inference of object-continuity in the face of dynamic input. That is, there is a single object representation, which initially contains information about the target, but masking occurs when it is modified to incorporate the mask. Altogether, as reviewed above, I propose that the evidence is overwhelmingly in favour of the object-updating framework for understanding OSM, and that it is time to lay to rest the notion of object-substitution. While the possibility that substitution may play a role in addition to object-updating cannot be ruled out, this is at best a modest role, and at present there is little evidence for even this. It is exciting to have this clarity about what OSM reflects because it means that it can be used as a scientific tool with greater precision and confidence. It also means that OSM belongs to the larger family of object-updating related phenomena including RB, the flash-lag effect, and apparent motion, but it has also made its own unique contribution to our understanding of the nature of updated representations (Harrison et al., 2016; Kahan & Enns, 2010), how and when updating occurs (Guest et al., 2012), and how updating is independent from visual attention (Filmer et al., 2015; Goodhew & Edwards, 2016).

The conclusion that OSM reflects object-updating also highlights how pervasive and powerful inferences of object-updating are in terms of our visual perception of the world around us (Enns et al., 2010). Without such forces, every change in location or appearance of an object would result in the inference that a new object had appeared in the scene. There would be no continuing objects. It is therefore clearly adaptive that such mechanisms exist. However, such forces need to be balanced against the requirement to be able to identify and recognise the onset of new objects when they do appear. The tension between these two goals – object continuity versus new object detection – is a fundamental challenge that our brain faces. The conclusion in favour of object-updating highlights that we can utilise OSM as a method to garner insight into the inferences and assumptions that the brain makes in the face of the challenge. While such inferences occur non-consciously, they ultimately shape our perception and understanding of the world around us.

For future research, it would be interesting to investigate questions such as how context and prior experience influence our expectancies and inferences about object continuity. Furthermore, how do different forms of psychopathology influence the tendency to individuate objects versus form persisting object representations? OSM can be used to answer such questions. Moreover, while preliminary evidence indicates that semantic information can influence the way in which object representations are assigned and updated (Hsu, Taylor, & Pratt, 2015), the vast majority of research on object-updating has examined the influence of surface features (e.g., colour), which do not necessarily violate the boundaries of object categories. For example, a green disc versus a blue disc – could still both be instances of the same disc, just changing colour or in different lighting conditions over time. While changes in such features degrade object-continuity inferences, they do not prevent them altogether. But what if the stimuli were such that the first and second instances could not possibly be the same object: such as an individual's face, followed by a house? Do the featural changes that cross object-category boundaries more severely undermine the process of object updating? Such knowledge will be important for 'scaling up' from well-controlled laboratory experiments with relatively impoverished stimuli (e.g., geometric shapes), to more 'real-world' contexts containing complex and recognisable objects – which is ultimately what we seek to understand.

Finally, visual-cognitive phenomena are often named according to their first discovered properties or proposed mechanisms. Case in point – the attentional blink, thought to be akin to actual eye-blink in suppressing all visual processing, is now known to permit very high-level processing of the second target, to the level of semantics (meaning), despite degraded conscious perception (Dux & Coltheart, 2005; Shapiro, Driver, Ward, & Sorensen, 1997; Vogel et al., 1998). So the comparison to an eye-blink now appears problematic given the level of suppression is so late in processing, but it is difficult to change once it has become popularised. So the shorthand ‘OSM’ is going to be difficult to change. But perhaps rather than standing for ‘object-*substitution* masking’, it could be agreed to refer to ‘object-*segmentation* masking’, where the segmentation refers to the process of creating either a continuing representation over time, versus segmenting or individuating distinct object representations over time.

Acknowledgements

This research was supported by an Australian Research Council (ARC) Discovery Early Career Research Award (DE140101734) awarded to S.C.G. Correspondence regarding this study should be addressed to Stephanie Goodhew (stephanie.goodhew@anu.edu.au), Research School of Psychology, The Australian National University.

References

Argyropoulos, I., Gellatly, A., Pilling, M., & Carter, W. (2013). Set size and mask duration do not interact in object-substitution masking. *Journal of Experimental Psychology: Human Perception and Performance, 39*(3), 646-661. doi: 10.1037/a0030240

Bavelier, D. (1994). Repetition blindness between visually different items: the case of pictures and words. *Cognition, 51*(3), 199-236. doi: 10.1016/0010-0277(94)90054-X

Binsted, G., Brownwell, K., Vorontsova, Z., Heath, M., & Saucier, D. (2007). Visuomotor system uses target features unavailable to conscious awareness. *Proceedings of the National Academy of Sciences, 104*(31), 12669-12672. doi: 10.1073/pnas.0702307104

Boehler, C. N., Schoenfeld, M. A., Heinze, H. J., & Hopf, J. M. (2008). Rapid recurrent processing gates awareness in primary visual cortex. *Proceedings of the National Academy of Sciences of the United States of America, 105*(25), 8742-8747. doi: 10.1073/pnas.0801999105

Bonneh, Y. S., Cooperman, A., & Sagi, D. (2001). Motion-induced blindness in normal observers. *Nature Neuroscience, 411*, 798-801.

Bouma, H. (1970). Interaction effects in parafoveal letter recognition. *Nature, 226*(5241), 177-178.

Bouvier, S., & Treisman, A. (2010). Visual feature binding requires reentry. *Psychological Science, 21*(2), 200-204. doi: 10.1177/0956797609357858

Breitmeyer, B. G., Herzog, M. H., & Ogmen, H. (2008). Motion, not masking, provides the medium for feature attribution. *Psychological Science, 19*(8), 823-829. doi: 10.1111/j.1467-9280.2008.02163.x

Breitmeyer, B. G., & Ögmen, H. (2000). Recent models and findings in visual backward masking: A comparison, review, and update. *Perception & Psychophysics, 62*(8), 1572-1595. doi: 10.3758/BF03212157

Breitmeyer, B. G., & Ögmen, H. (2006). *Visual masking: Time slices through conscious and unconscious vision* (2nd ed.). New York, NY: Oxford University Press.

Burke, L. (1952). On the tunnel effect. *Quarterly Journal of Experimental Psychology, 4*(3), 121-138. doi: 10.1080/17470215208416611

Camp, S. J., Pilling, M., Argyropoulos, I., & Gellatly, A. (2015). The Role of Distractors in Object Substitution Masking. *Journal of Experimental Psychology: Human Perception and Performance*. doi: 10.1037/xhp0000065 ·

Cartwright-Finch, U., & Lavie, N. (2007). The role of perceptual load in inattentional blindness. *Cognition, 102*(3), 321-340. doi: 10.1016/j.cognition.2006.01.002

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21*, 109-127.

Dell'Acqua, R., Dux, P. E., Wyble, B., Doro, M., Sessa, P., Meconi, F., & Jolicoeur, P. (2015). The attentional blink impairs detection and delays encoding of visual information: evidence from human electrophysiology. *Journal of Cognitive Neuroscience, 27*(4), 720-735. doi: 10.1162/jocn_a_00752

Dell'Acqua, R., Dux, P. E., Wyble, B., & Jolicoeur, P. (2012). Sparing from the attentional blink is not spared from structural limitations. *Psychonomic Bulletin & Review, 19*(2), 232-238. doi: 10.3758/s13423-011-0209-3

Di Lollo, V. (2010). Iterative reentrant processing: A conceptual framework for perception and cognition (the blinding problem? No worries, mate). In V. Coltheart (Ed.), *Tutorials in Visual Cognition* (pp. 9-42). New York: Psychology Press.

Di Lollo, V., Enns, J. T., & Rensink, R. A. (2000). Competition for consciousness among visual events: The psychophysics of reentrant visual processes. *Journal of*

Experimental Psychology: General, 129(4), 481-507. doi: 10.1037/0096-3445.129.4.481

Dux, P. E., & Coltheart, V. (2005). The meaning of the mask matters: Evidence of conceptual interference in the attentional blink. *Psychological Science, 16,* 775-779.

Dux, P. E., & Coltheart, V. (2008). Repetition blindness and repetition priming: Effects of featural differences between targets and distractors on RSVP dual-target search. *Memory & Cognition, 36*(4), 776-790. doi: 10.3758/MC.36.4.776

Dux, P. E., & Marois, R. (2009). How humans search for targets through time: A review of data and theory from the attentional blink. *Attention, Perception, & Psychophysics, 71*(8), 1683-1700. doi: 10.3758/APP.71.8.1683

Eagleman, D. M., & Sejnowski, T. J. (2000). Motion integration and postdiction in visual awareness. *Science, 287*(5460), 2036-2038.

Enns, J. T. (2002). Visual binding in the standing wave illusion. *Psychonomic Bulletin & Review, 9*(3), 489-496.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science, 8*(2), 135-139. doi: 10.1111/j.1467-9280.1997.tb00696.x

Enns, J. T., & Di Lollo, V. (2000). What's new in visual masking? *Trends in Cognitive Sciences, 4*(9), 345-352. doi: 10.1016/S1364-6613%2800%2901520-5

Enns, J. T., Lleras, A., & Moore, C. M. (2010). Object updating: A force for perceptual continuity and scene stability in human vision. In R. Nijhawan & B. Khurana (Eds.), *Space and Time in Perception and Action* (pp. 503-520). Cambridge: Cambridge University Press.

Eriksen, C. W., & St. James, J. D. (1986). Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics, 40*(4), 225-240. doi: 10.3758/BF03211502

Filmer, H. L., Mattingley, J. B., & Dux, P. E. (2014). Size (mostly) doesn't matter: The role of set size in object substitution masking. *Attention, Perception, & Psychophysics*. doi: 10.3758/s13414-014-0692-5

Filmer, H. L., Mattingley, J. B., & Dux, P. E. (2015). Object substitution masking for an attended and foveated target. *Journal of Experimental Psychology: Human Perception and Performance*. doi: 10.1037/xhp0000024

Flombaum, J. I., Kundey, S. M., Santos, L. R., & Scholl, B. J. (2004). Dynamic object individuation in rhesus macaques. *Psychological Science, 15*(12), 795-800.

Flombaum, J. I., & Scholl, B. J. (2006). A temporal same-object advantage in the tunnel effect: Facilitated change detection for persisting objects. *Journal of Experimental Psychology: Human Perception and Performance, 32*(4), 840-853. doi: 10.1037/0096-1523.32.4.840

Gellatly, A., Pilling, M., Carter, W., & Guest, D. (2010). How does target duration affect object substitution masking. *Journal of Experimental Psychology: Human Perception and Performance, 36*(5), 1267-1279. doi: 10.1037/a0018733

Giesbrecht, B., Bischof, W. F., & Kingstone, A. (2003). Visual masking during the attentional blink: tests of the object substitution hypothesis. *Journal of Experimental Psychology: Human Perception and Performance*, *29*(1), 238-258.

Goldfarb, L., & Treisman, A. (2011). Repetition blindness: The survival of the grouped. *Psychonomic Bulletin & Review, 18*(6), 1042-1049. doi: 10.3758/s13423-011-0135-4

Goodhew, S. C. (2016). When Masks Reveal More Than They Hide. *The American Journal of Psychology, 129*(3), 350-355. doi: 10.5406/amerjpsyc.129.3.0350

Goodhew, S. C., Dux, P. E., Lipp, O. V., & Visser, T., A. W. (2012). Understanding recovery from object substitution masking. *Cognition, 122*(3), 405-415. doi: 10.1016/j.cognition.2011.11.010

Goodhew, S. C., & Edwards, M. (2016). Object individuation is invariant to attentional diffusion: Changes in the size of the attended region do not interact with object-substitution masking. *Cognition, 157*, 358-364. doi: 10.1016/j.cognition.2016.10.006

Goodhew, S. C., Edwards, M., Boal, H. L., & Bell, J. (2015). Two objects or one? Similarity rather than complexity determines objecthood when resolving dynamic input. *Journal of Experimental Psychology: Human Perception and Performance, 41*(1), 102-110. doi: 10.1037/xhp0000022

Goodhew, S. C., Fogel, N., & Pratt, J. (2014). The nature of altered vision near the hands: Evidence for the magnocellular enhancement account from object correspondence through occlusion. *Psychonomic Bulletin & Review, 21*(6), 1452-1458. doi: 10.3758/s13423-014-0622-5

Goodhew, S. C., Greenwood, J. A., & Edwards, M. (2016). Categorical information influences conscious perception: An interaction between object-substitution masking and repetition blindness. *Attention Perception & Psychophysics, 78*(4), 1186-1202. doi: 10.3758/s13414-016-1073-z

Goodhew, S. C., Pratt, J., Dux, P. E., & Ferber, S. (2013). Substituting objects from consciousness: A review of object substitution masking. *Psychonomic Bulletin & Review, 20*(5), 859-877. doi: 10.3758/s13423-013-0400-9

Goodhew, S. C., Shen, E., & Edwards, M. (2016). Selective spatial enhancement: Attentional spotlight size impacts spatial but not temporal perception. *Psychonomic Bulletin & Review, 23*(4), 1144-1149. doi: 10.3758/s13423-015-0904-6

Goodhew, S. C., Visser, T., A. W., Lipp, O. V., & Dux, P. E. (2011). Implicit semantic perception in object substitution masking. *Cognition, 118*(1), 133-137. doi: 10.1016/j.cognition.2010.10.013

Green, M. F., Wynn, J., Breitmeyer, B., Mathis, K., & Nuechterlein, K. (2011). Visual masking by object substitution in schizophrenia. *Psychological Medicine, 41*(7), 1489-1496. doi: 10.1017/S003329171000214X

Greenwood, J. A., Bex, P. J., & Dakin, S. C. (2010). Crowding changes appearance. *Current Biology, 20*(6), 496-501.

Guest, D., Gellatly, A., & Pilling, M. (2011). The effect of spatial competition between object-level representations of target and mask on object substitution masking. *Attention, Perception, and Psychophysics, 73*(8), 2528-2541. doi: 10.3758/s13414-011-0196-5

Guest, D., Gellatly, A., & Pilling, M. (2012). Reduced OSM for long duration targets: Individuation or items loaded into VSTM? *Journal of Experimental Psychology: Human Perception and Performance, 38*(6), 1541-1553. doi: 10.1037/a0027031

Harrison, G. W., Rajsic, J., & Wilson, D. E. (2016). Object-substitution masking degrades the quality of conscious object representations. *Psychonomic Bulletin & Review, 23*(1), 180-186. doi: 10.3758/s13423-015-0875-7

Hein, E., & Moore, C. M. (2012). Spatio-temporal priority revisited: The role of feature identity and similarity for object correspondence in apparent motion. *Journal of Experimental Psychology: Human Perception and Performance, 38*(4), 975-988. doi: 10.1037/a0028197

Herzog, M. H., & Koch, C. (2001). Seeing properties of an invisible object: Feature inheritance and shine-through. *Proceedings of the National Academy of Sciences of the United States of America, 98,* 4271-4275. doi: 10.1073/pnas.071047498

Hirose, N., Kihara, K., Mima, T., Ueki, Y., Fukuyama, H., & Osaka, N. (2007). Recovery from object substitution masking induced by transient suppression of visual motion processing: A repetitive transcranial magnetic stimulation study. *Journal of Experimental Psychology: Human Perception and Performance, 33*(6), 1495-1503. doi: 10.1037/0096-1523.33.6.1495

Hollingworth, A., & Franconeri, S. L. (2009). Object correspondence across brief occlusion is established on the basis of both spatiotemporal and surface feature cues. *Cognition, 113*(2), 150-166. doi: 10.1016/j.cognition.2009.08.004

Hsu, P., Taylor, J. E., & Pratt, J. (2015). Frogs Jump Forward: Semantic Knowledge Influences the Perception of Element Motion in the Ternus Display. *Perception, 44*(7), 779-789.

Ibbotson, M., & Krekelberg, B. (2011). Visual perception and saccadic eye movements. *Current Opinion in Neurobiology, 21*(4), 553-558. doi: 10.1016/j.conb.2011.05.012

Jahshan, C., Wynn, J. K., McCleery, A., Glahn, D. C., Altshuler, L. L., & Green, M. F. (2014). Cross-diagnostic comparison of visual processing in bipolar disorder and schizophrenia. *Journal of psychiatric research, 51*, 42-48. doi: 10.1016/j.jpsychires.2013.12.014

Kahan, T. A., & Enns, J. T. (2010). Object trimming: When masking dots alter rather than replace target representations. *Journal of Experimental Psychology: Human Perception and Performance, 36*(1), 88-102. doi: 10.1037/a0016466

Kahan, T. A., & Lichtman, A. S. (2006). Looking at object-substitution masking in depth and motion: Toward a two-object theory of object substitution. *Perception & Psychophysics, 68*(3), 437-446. doi: 10.3758/BF03193688

Kahneman, D., & Treisman, A. (1984). Changing views of attention and automaticity. In R. Parasuraman & D. A. Davies (Eds.), *Varieties of attention*. New York: Academic Press.

Kahneman, D., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, *24*(2), 175-219. doi: 10.1016/0010-0285%2892%2990007-O

Kanwisher, N. (1987). Repetition blindness: Type recogniton without token individuation. *Cognition, 27*(2), 117-143. doi: 10.1016/0010-0277%2887%2990016-3

Kanwisher, N. (1991). Repetition blindness and illusory conjunctions: Errors in binding visual types with visual tokens. *Journal of Experimental Psychology: Human Perception and Performance, 17*(2), 404-421. doi: 10.1037/0096-1523.17.2.404

Kanwisher, N., & Driver, J. (1992). Objects, attributes, and visual attention: Which, what and where. *Current Directions in Psychological Science, 1*(1), 26-31. doi: 10.1111/1467-8721.ep10767835

Kanwisher, N., Kim, J. W., & Wickens, T. D. (1996). Signal detection analyses of repetition blindness. *Journal of Experimental Psychology: Human Perception and Performance, 22*, 1249-1260.

Kanwisher, N., & Potter, M. C. (1989). Repetition blindness: The effects of stimulus modality and spatial displacement. *Memory & Cognition, 17*(2), 117-124. doi: 10.3758/BF03197061

Kim, C. Y., & Blake, R. (2005). Psychophysical magic: Rendering the visible 'invisible'. *Trends in Cognitive Sciences, 9*(8), 381-388. doi: 10.1016/j.tics.2005.06.012

Kotsoni, E., Csibra, G., Mareschal, D., & Johnson, M. H. (2007). Electrophysiological correlates of common-onset visual masking. *Neuropsychologia, 45*(10), 2285-2293. doi: 10.1016/j.neuropsychologia.2007.02.023

LaBerge, D. (1983). Spatial extent of attention to letters and words. *Journal of Experimental Psychology: Human Perception and Performance, 9*(3), 371-379. doi: 10.1037/0096-1523.9.3.371

Lamme, V. A. F., & Roelfsema, P. R. (2000). The distinct modes of vision offered by feedforward and recurrent processing. *Trends in Neurosciences, 23*(11), 571-579. doi: 10.1016/S0166-2236(00)01657-X

Lee, J., Cohen, M. S., Engel, S. A., Glahn, D., Nuechterlein, K. H., Wynn, J. K., & Green, M. F. (2014). Neural substrates of visual masking by object substitution in schizophrenia. *Hum Brain Mapp, 35*(9), 4654-4662. doi: 10.1002/hbm.22501

Lim, S. W. H., & Chua, F. K. (2008). Object substitution masking: When does mask preview work? *Journal of Experimental Psychology: Human Perception and Performance, 34*(5), 1108-1115. doi: 10.1037/0096-1523.34.5.1108

Lleras, A., & Moore, C. M. (2003). When the target becomes the mask: Using apparent motion to isolate the object-level component of object substitution masking. *Journal of Experimental Psychology: Human Perception and Performance, 29*(1), 106-120. doi: 10.1037/0096-1523.29.1.106

Luiga, I., & Bachmann, T. (2008). Luminance processing in object substitution masking. *Vision Research, 48*(7), 937-945. doi: 10.1016/j.visres.2008.01.001

Luiga, I., Gellatly, A., & Bachmann, T. (2010). Delayed offset of distractors masks a local target. *Acta Psychologica, 134*(3), 344-352.

Matin, E. (1974). Saccadic suppression: A review and an analysis. *Psychological Bulletin, 81*(12), 899-917. doi: 10.1037/h0037368

Mikami, A. (1991). Direction selective neurons respond to short-range and long-range apparent motion stimuli in macaque visual area MT. *Int J Neurosci, 61*(1-2), 101-112.

Moore, C. M., & Enns, J. T. (2004). Object Updating and the Flash-Lag Effect. *Psychological Science, 15*(12), 866-871. doi: 10.1111/j.0956-7976.2004.00768.x

Moore, C. M., & Lleras, A. (2005). On the role of object representations in substitution masking. *Journal of Experimental Psychology: Human Perception and Performance, 31*(6), 1171-1180. doi: 10.1037/0096-1523.31.6.1171

Moore, C. M., Stephens, T., & Hein, E. (2010). Features, as well as space and time, guide object persistence. *Psychonomic Bulletin & Review, 17*(5), 731-736. doi: 10.3758/PBR.17.5.731

Most, S. B., Scholl, B. J., Clifford, E. R., & Simons, D. J. (2005). What you see is what you set: sustained inattentional blindness and the capture of awareness. *Psychological Review, 112*(1), 217-242. doi: 10.1037/0033-295x.112.1.217

Muckli, L., Kohler, A., Kriegeskorte, N., & Singer, W. (2005). Primary Visual Cortex Activity along the Apparent-Motion Trace Reflects Illusory Perception. *PLoS Biol, 3*(8), e265. doi: 10.1371/journal.pbio.0030265

Muckli, L., Kriegeskorte, N., Lanfermann, H., Zanella, F. E., Singer, W., & Goebel, R. (2002). Apparent motion: event-related functional magnetic resonance imaging of perceptual switches and States. *The Journal of neuroscience : the official journal of the Society for Neuroscience, 22*(9), RC219. doi: 20026362

Muller, N. G., Bartelt, O. A., Donner, T. H., Villringer, A., & Brandt, S. A. (2003). A physiological correlate of the "zoom lens" of visual attention. *Journal of Neuroscience, 23*(9), 3561-3565.

Neill, W. T., Hutchison, K. A., & Graves, D. F. (2002). Masking by object substitution: Dissociation of masking and cueing effects. *Journal of Experimental Psychology: Human Perception and Performance, 28*(3), 682-694. doi: 10.1037/0096-1523.28.3.682

Nijhawan, R. (2002). Neural delays, visual motion and the flash-lag effect. *Trends in Cognitive Sciences, 6*(9), 387-393. doi: 10.1016/S1364-6613(02)01963-0

Pelli, D. G., & Tillman, K. A. (2008). The uncrowded window of object recognition. *Nature Neuroscience, 11*(10), 1129-1135. doi: 10.1038/nn.2187

Pilling, M., & Gellatly, A. (2010). Object substitution masking and the object updating hypothesis. *Psychonomic Bulletin & Review, 17*(5), 737-742. doi: 10.3758/PBR.17.5.737

Raymond, J. E., Shapiro, K. L., & Arnell, K. M. (1992). Temporary suppression of visual processing in an RSVP task: An attentional blink? *Journal of Experimental Psychology: Human Perception and Performance, 18*(3), 849-860.

Richard, A. M., Luck, S. J., & Hollingworth, A. (2008). Establishing object correspondence across eye movements: Flexible use of spatiotemporal and surface feature information. *Cognition, 109*(1), 66-88. doi: 10.1016/j.cognition.2008.07.004

Rock, I., Linnet, C. M., Grant, P., & Mack, A. (1992). Perception without attention: Results of a new method. *Cognitive Psychology, 24*(4), 502-534. doi: 10.1016/0010-0285(92)90017-V

Shapiro, K., Driver, J., Ward, R., & Sorensen, R. E. (1997). Priming from the attentional blink: A failure to extract visual tokens but not visual types. *Psychological Science, 8*, 95-100.

Tsal, Y., & Benoni, H. (2010). Diluting the burden of load: perceptual load effects are simply dilution effects. *Journal of Experimental Psychology: Human Perception and Performance, 36*(6), 1645-1656. doi: 10.1037/a0018172

Vogel, E. K., Luck, S. J., & Shapiro, K. L. (1998). Electrophysiological evidence for a postperceptual locus of suppression during the attentional blink. *Journal of Experimental Psychology: Human Perception and Performance, 24*, 1656-1674.

Wallis, T. S. A., & Arnold, D. H. (2009). Motion-Induced Blindness and Motion Streak Suppression. *Current Biology, 19*(4), 325-329. doi: http://dx.doi.org/10.1016/j.cub.2008.12.053

Yi, D. J., Turk-Browne, N. B., Flombaum, J. K., Kim, M. S., Scholl, B. J., & Chun, M. M. (2008). Spatiotemporal object continuity in human ventral visual cortex. *Proceedings of the National Academy of Sciences, 105*(26), 8840-8845. doi: 10.1073/pnas.0802525105