THE EPISODIC NATURE OF CONSCIOUSNESS: AN INVESTIGATION USING RAPID SERIAL VISUAL PRESENTATION

By

Alberto Avilés

A thesis submitted to the University of Birmingham for the degree of DOCTOR OF PHILOSOPHY

School of Psychology

College of Life and

Environmental Sciences

November 2022

ABSTRACT

What is consciousness for? Which cognitive functions can only be carried out with consciousness? Although much empirical work has explored these questions there is still not a definitive answer. In this thesis, we review how the evidence obtained with the Rapid Serial Visual Presentation(RSVP) paradigm could shed light on this issue. A key finding is the low recognition performance for RSVP stimuli that contrasts with the excellent performance in detection and identification tasks. This is suggestive of mechanisms that are highly efficient in the absence of conscious reports. Based on this effect and some of the theoretical ideas accounting for key RSVP effects we hypothesize that one cognitive mechanism that can only be carried out with consciousness is episodic encoding.

The thesis presents a series of experiment testing this hypothesis. Classic episodic memory effects are explored in the context of RSVP experiments. We address the repetition effect, the expected improvement in memory consolidation as items are repeated. Recency effects are investigated in the contexts of recognition and free-recall tasks. Proactive interference and list-length effect are also addressed using RSVP stimuli. Finally, we tested the reportability of task-irrelevant items.

Our main aim is to provide a comprehensive picture of the episodic memory effects that can be observed in RSVP. The results show that many of our findings are consistent with the view that most of the stimuli presented in RSVP sequences are not episodically encoded. This challenges the possibility of unconscious episodic encoding and agrees with our main hypothesis. These findings are discussed in the context of a theoretical proposal about the function of consciousness. The Tokenized Percept Hypothesis suggests that the unconscious activation of categorical or featural information might be enough for processes, such as detection or identification. While the acquisition of episodic properties necessitates conscious processing.

ACNOWLEDGEMENTS

I have used 38098 words to present this project, but many more would be needed to convey my gratitude to the people that made all this possible. First and foremost, I thank my principal supervisor, Professor Howard Bowman. I feel privileged for the chance to work with you. After quite a few years and countless meetings, it keeps being a joyful experience to learn from you. I will forever be indebted to you for the wisdom, support and patience that you have shared with me. Especially, I feel grateful for the opportunity of exploring and developing your ideas about consciousness. All this time thinking about consciousness, time, types and tokens has been pure joy and excitement. Thanks for having more confidence in me than myself. And thank you for offering me one of the rarest and most precious gifts in life: a second chance.

I would like to say thank you to Professors Jane Raymond, Kim Shapiro and Howard Bowman for opening the doors of their lab to me. Literally, this work would not have been possible without their support. I have grown as a researcher with every lab meeting. It has been a great privilege to be part of it. A huge thanks to Pia Rothstein for your very insightful comment and the magnificent breakfast meetings. I am grateful to all the fellow researchers in the group who always offered kind words and a helpful hand. Especially, I want to thank Sara Assecondi for all her efforts to make me feel part of the group. Your support has been an essential part of the progress I have made during my PhD. Thanks to Dr Daniel Dogson and Dr Kelly Garner for your kindness and wisdom. Aleena Mahmood, Claire Thirkettle, Jenna Haxby and all the project students, thank you for your enthusiasm and your help in recruiting and testing participants.

Many thanks to my friends in Oxford. Fidel, Ruben, Dan, and Nayeli, without you this would have been insurmountable. Thanks for the fun, and thanks for pushing me.

To Olivia, thanks for your love, your patience -and your patience again. From the very first day to the last of this long adventure you have pushed, consoled, helped, and encouraged me. We have crossed this bridge together.

To my Mum. Thank you for being there and for your unwavering belief in me.

TABLE OF CONTENTS

Chapter 1: General Introduction

> There are several reasons why consciousness is so baffling. For one thing, it seems to be among the chronically unemployed… The question then arises: what does consciousness add to what unconsciousness can achieve? …. As far as anybody knows, anything that our conscious minds do they could do just as well if they weren't conscious. Why, then, did God bother to make consciousness? What on earth could he have had in mind?
>
> Jerry Fodor (2004)

1. The function of consciousness

This thesis presents an investigation into the nature of our consciousness, which, from an access awareness perspective, corresponds to our reportable subjective experience (Dehaene & Changeux, 2011). More concretely, this is an investigation about the function of consciousness. This issue has been approached from, at least, three different perspectives. First, consciousness might be an epiphenomenon, a by-product of some biological function, that does not play any functional role. Second, consciousness may have a biological function that resulted in traits favoured by natural selection. According to this view, consciousness provides an advantage to the agent (Cleeremans & Tallon-Baudry, 2022; Frith & Metzinger, 2016). The third view emphasizes that one needs to find first what consciousness does in order to know what consciousness is for. According to this view, consciousness is assimilated into the functional language of cognitive sciences and the goal is to discover what is the function of consciousness.

Regardless of the point of view that one takes about this issue, it is clear that any compelling solution would be valuable not only to the philosophical discussion about the hard problem of consciousness (the challenge of understanding how and why subjective conscious experiences arise from physical processes in the brain) (Chalmers, 1996)[1], but also to accommodate consciousness into the scientific view of the world. For example, to fit consciousness into the existing knowledge of the evolution of cognition requires researchers to determine what biological function it fulfils. More importantly, for the purposes of the present thesis, to know what consciousness does would allow researchers to insert the study of consciousness into the realm of cognitive sciences in general and psychology in particular.

If a functional account of consciousness is successful, one would be tempted to conclude that consciousness is nothing more than a functional process. Once the functional properties are determined, we would know all there is to know about consciousness as is the case of other cognitive processes. To know what Attention *does* is to know what Attention *is.* Consciousness might be a cognitive process and that's that. Unfortunately, there is no obvious place for consciousness in the functional language of psychology. In the study programs for a Degree in Psychology, consciousness is rarely found as a subject.

[1] In fact one of the core debates in the Hard Problem of Consciousness concerns the conceivability of Philosophical Zombies, an agent that is physically and *functionally* identical to a normal person but without the qualitative aspect of consciousness. Both the concept of the Hard Problem of consciousness and the notion of Philosophical Zombies have engendered substantial controversy, eliciting extensive debate and challenging their respective contributions to the advancement of scientific comprehension of consciousness.

The lack of an obvious fit of consciousness among the major cognitive processes has forced researchers to approach the issue of the function of consciousness from a different angle. In a similar research endeavour to the quest for the Neural Correlates of Consciousness (NCC), perhaps a good first step is to search for the Functional Correlates of Consciousness (FCC). An NCC can be defined as the minimal neuronal mechanisms jointly sufficient for any one specific conscious percept (Tononi & Koch, 2008). Similarly, an FCC is a function that can only be carried out with consciousness. By function we mean one of the cognitive functions, "those processes by which the sensory input is transformed, reduced, elaborated, stored, recovered, and used"(Neisser, 2014, p. 4).

2. The Functional Correlates of Consciousness

Executive functions

One popular conception of consciousness situates consciousness among the high-level cognitive processes. Some researchers think of consciousness as the cusp of cognition, and in this view many of the high-level cognitive functions are accompanied by consciousness. Particularly relevant are executive functions such as control or monitoring. For example, Shallice (1972) proposed that consciousness would be related with the function of setting a goal for the dominant action plan; and Norman and Shallice (1986) associated consciousness with the function of cognitive control. Cognitive control is involved in switching between tasks or interrupting a dominant task. However, even for this high-level mechanism, there is some evidence that suggest that these cognitive functions can be carried out in the absence of consciousness. Lau and Passingham (2007) presented an experiment where participants had to switch tasks (phonological task or semantic) depending on a cue. They showed that when the cue for one task was preceded by a subliminally presented cue for the other task, participants were less accurate and slower. Similarly, van Gaal et al. (2010) focused on another higher-order process: response inhibition. In a go/no-go task, that included some non-conscious no-go trials, they showed that unconscious no-go stimulus resulted in a significant slowdown in the speed of responding (in fact, this effect is quite small, 10 ms, relative to the much larger differences observed in the "conscious" condition. This might be interpreted as supporting a view where there are differences associated with cognitive controls between conscious and non-conscious conditions). Similar evidence has resulted in the questioning of whether consciousness is really associated with cognitive control (Hassin, 2013). Therefore, cognitive control might not be a FCC after all.

Working memory (WM)

WM can be defined as "a limited capacity system allowing the temporary storage and manipulation of information necessary for such complex tasks as comprehension, learning and reasoning " (Baddeley, 2000, p. 418). The most influential model of WM is Baddeley and Hitch's WM model (1974). This model consists of a series of buffers: the visuospatial sketchpad, the phonological buffer, and an attentional control system, the "central executive", with the function of establishing goal-directed behaviour. This model was updated by Baddeley (2000) to incorporate a crucial component for consciousness, the episodic buffer. This is a limited capacity store to form episodes through the binding of multimodal information (Baddeley, 2003). It is controlled by the central executive, and it stores conscious representations. In fact, it has been suggested that the episodic buffer could store representations that could be globally broadcasted to other cognitive systems. This function attributed to the episodic buffer is one of the main connections between Baddeley's Working Memory model and one of the leading theories of consciousness, the Global Workspace Theory (Baars, 1993; Baars & Franklin, 2003).

The concept of Working Memory has evolved to integrate neuropsychological findings. It has been suggested (Postle, 2006) that the retention of information in WM is associated with sustained activity in the same brain regions responsible for representing that information in non-working memory situations, such as perception, semantic memory, motor control, and language processing. Secondly, humans naturally employ multiple mental codes when representing stimuli in WM, opportunistically utilizing as many codes as are available. This enables the incorporation of various aspects related to the stimulus, including previous experiences, contextual information, and emotional associations. These principles align with the notion that WM involves the temporary activation of representations in long-term memory.

The concept of Working Memory is closely related to the earlier concept of Short-term memory (Atkinson & Shiffrin, 1971). This is a store that holds information temporarily in the absence of perceptual stimulation. Most of the evidence exploring the relationship between WM and consciousness focuses on the aspects of WM that are shared with the concept of Short-term memory. For example, Soto et al. (2011) presented participants with masked orientation Gabor cues. After a delay period (2 s or 5 s in different experiments) participants were asked to respond to a Gabor orientation test. In addition, participants were asked to evaluate from 1 ("did not see anything") to 4 ("saw the stimulus and its orientation") the awareness of the memory cue. They found above chance performance (accuracy above 50%) in the Gabor test even when only the trials rated as 1 in the awareness scale were included. They concluded that "visual memory can encode, maintain and access unconscious items" (Soto et al., 2011, p. 51). This challenged the notion that WM or, visual Short-term memory can only store representations that have been consciously perceived. In an experiment testing a similar issue, Trübutschek et al. (2017) obtained results consistent with the view that non-conscious representations could be retrieved after a delay period and they can be stored as a form of “activity silent” (see Stokes, 2015) WM . In conclusion, recent evidence seems to indicate that visual Short-term memory can hold non-conscious information that can be accessed after several seconds . These findings represents a challenge for the view that WM is a functional correlate of consciousness.

Attention

We will use the term attention to refer to the prioritization of some information at the expense of other. Perhaps one of the models of attention that more often has been discussed regarding the relationship between attention and consciousness is the biased competition model (Desimone & Duncan, 1995). This model posits that the visual input of objects compete for limited processing resources. Attended stimuli are those that win the competition and both, top-down and bottom-up influences, can bias the competition.

The ideas put forward to account for the relationship between attention and consciousness go beyond the claim that the two processes are identical. Instead, the discussion is about the extent to which attention is necessary for consciousness, or the opposite, whether there could be consciousness without attention. There is some evidence showing that attention can be directed to unconscious stimuli and amplify their behavioural impact. For example, Naccache et al. (1997) showed that unconscious priming depended on the allocation of attention to the prime-target pair. They designed a masked priming task where participants had to indicate if a target number was bigger or smaller than the number 5. They manipulated the congruency of the primes along with the predictability of the position of the prime-target pair in a sequence of rapidly presented stimuli. They found facilitation (better performance in congruent prime-target pairs) for unconscious cues in predictable sequences. This finding is consistent with attention being directed to unconscious stimuli. Similarly, Jiang et al. (2006) using Continuous Flash Suppression (CFS) (Tsuchiya & Koch, 2005) showed that invisible erotic stimuli can impact the distribution of spatial attention. CFS is a relatively recent paradigm that consists of the presentation of a series of masks to the dominant eye concurrently to the presentation of the target information to the non-dominant eye. They presented erotic pictures rendered invisible with CFS to the left or the right part of the screen. Participants were asked to respond to a later Gabor test presented either to the left or the right. The results showed that reaction times were modulated by the consistency (same side of the screen) between the pictures and the Gabor cue, and the gender and sexual orientation of the participants. In agreement with the view that consciousness is not necessary for attention, Kentridge et al., (1999) showed that there is attention to stimuli presented to the blind hemifield of blindsighters.

Although the evidence showing that it is possible to deploy attention to nonconscious stimuli is compelling, the opposite claim of conscious perception without attention has not been supported by the literature (Cohen & Chun, 2017). Perhaps, the most striking evidence is the Inattention Blindness effect (Mack & Rock, 1998), the failure to report clearly visible but unattended stimuli. For example, participants might fail to report a clearly visible object when it is presented alongside a task-relevant stimulus. Mack & Rock (1998 p.71) interpreted this effect as supporting the view that "attention provides the key that unlocks the gate dividing unconscious perception (which, according to our working hypothesis, entails deep processing) from conscious perception".

Interestingly, it has been found that consciousness alters the appearance of stimuli (Carrasco et al., 2004). They showed that attention boosts the apparent stimulus contrast. This suggests that attention can enhance the saliency of stimuli by increasing their contrasts. This is consistent with the view that attention and consciousness are closely related.

Episodic memory

Perhaps the cognitive function that we intuitively associate more closely with consciousness is episodic memory. It is difficult to conceive that we could recall the what-where-when of an event without being aware of the corresponding memory. In fact, the close relationship between consciousness and episodic memory is more than a mere intuition and it has been part of the study of memory since the origins of cognitive neuropsychology. The study of amnesic patients inspired the notion that consciousness could play a role as a criterion to differentiate between memory systems (Squire & Zola, 1996). Memories could be differentiated between those that can be declared, declarative memory, and those that do not, nondeclarative memory. Declarative memory is assumed to be supported by conscious representations, and it consists of semantic and episodic memory. Nondeclarative memory would involve phenomena such as priming, skill learning or classical conditioning.

Another division of memory where consciousness played a role as a criterion is the division between explicit and implicit memory (Graf & Schacter, 1985). Explicit memory requires conscious recollection of the study episode, while implicit memory occurs when performance in a memory test is facilitated without consciously remembering. Implicit memory can be tested with word-fragment completion or lexical decision tasks, and is based on a priming effect, which facilitates performance. Explicit memory is assumed to be tested with recognition or free-recall tests.

However, other criteria have been suggested as more relevant than consciousness to differentiate memory systems. In the relational memory theory (Cohen et al., 1997), declarative memory is not ultimately defined by its relationship with consciousness. Instead, declarative memory is identified with a hippocampal system that captures arbitrary or accidental relations between features of an event or between different events. Henke (2010) agrees with the view that consciousness is a poor criterion to discern these memory systems. Instead, they proposed a processing-based model that depends on the speed of the encoding (rapid versus slow), nature of the memory representation (associative versus single-items), and the flexibility of the representations. Consciousness is not assumed to be especially relevant in any of these processing modes. Therefore, even processing features traditionally postulated as requiring conscious representations, such as the encoding of associations, could be achieved, in principle, with nonconscious representations. There is some evidence supporting exactly that. Reber & Henke (2011) reported the results of a memory experiment where participants were presented with pairs of masked words for 17 ms in the encoding phase. After a 5-min break, participants were tested with conceptually related words to the encoding material. They manipulated whether the words at test were an analogue (a pair of words conceptually related to a pair used at encoding) or a broken analogue (the corresponding words at encoding were not presented simultaneously). For example, the words "desk-bus", presented in the test phase are a good analogue of the pair "table-car", presented in the encoding phase ("desk-sheep" would be a broken analogue). Participants were asked whether the words at test fitted together. They found that the proportion of fitted responses was higher for "analogues" relative to the "broken analogues". These results support the notion that association can be encoded even when stimuli are subliminally presented. Other experiments have challenged the idea that consciousness is required for recognition or free-recall tasks. Chong et al., (2014) reported a series of experiments to test whether a subliminally presented word at encoding or test could result in above chance performance in a recognition task. For example, in their Experiment 2, they presented forward and backward masked words during the study phase and participants were tested with supraliminal cues. They observed a small (d' = 0.27) but significant above chance performance.

In summary, consciousness is no longer the favourite criterion used by theorists to classify memory systems. In addition, some of the memory processing features commonly associated with consciousness (e.g. rapid formation of association) have been observed in experiments using subliminal presentation of stimuli. Based on this, is episodic memory still a candidate for FCC? The jury is still out. The defining property of episodic memory is that it enables the retrieval of the what, where, and when of past events (Tulving, 1985, 2001). We concur with Reber & Henke (2011, p. 9) in the acknowledgment that "a firm claim of an unconscious form of episodic memory must await evidence of the unconscious formation of what-where-when memories".

Is consciousness a capacity booster?

From our very brief review, we can conclude that there is a considerable amount of evidence supporting the idea that there are no cognitive functions that are uniquely conscious. This evidence is so compelling that some researchers defend the slogan/principle "Yes It Can" (the unconscious can do the same things the conscious does) (Hassin, 2013; see also Rosenthal, 2008). However, this is still a highly debated topic. For example, Bauemeister and Masicampo (2010) reviewed empirical evidence concerning the functional contribution of conscious processed. In summary, the showed that consciousness is necessary for the construction of meaningful, sequential thought, and this function, sequential thought, would be essential to facilitate social and cultural interactions.

Others have argued that even for tasks where consciousness improves performance "it would be odd – within a broadly computational/cognitive approach to the mind - if it were absolutely impossible to do the same task non-consciously" (Shea & Frith, 2016, p. 3). Therefore, it has been argued that instead of searching for a uniquely conscious cognitive function, it could be more productive to think of consciousness as playing a facilitatory role for some processes and task (Shea & Frith, 2016). Particularly, consciousness could facilitate the integration of information in the brain. The mere notion that consciousness involves *reportable* experiences implies integration of information; since conscious information must be available to a range of consuming systems, such as, verbal reports or goal-directed reasoning. Is integration an inherently conscious function? Shea & Frith, (2016) argue that integration can be achieved on the basis of nonconscious representations, but it is facilitated by consciousness. Faivre et al., (2014) showed that there is integration of multimodal information from subliminal stimuli, but only after a training phase with supraliminal stimuli. Interestingly, this is also the case for the Reder & Henke's experiment (2011) described above. Thus, perhaps, the facilitatory role of consciousness applies to the encoding of episodic information too.

3. THEORIES OF CONSCIOUSNESS

Higher-order Theories of Consciousness

There are two main principles that define all Higher-order Theories (Brown et al., 2019). The first principle is the claim that first-order representations not targeted by higher-order processes are not sufficient for consciousness. Second, a state is conscious only when a subject is conscious of that state noninferentially and nonobservationally (that is, that state is not accessed, for example, by inferences drawn by observing your own behaviour) . This is known as the Transitivity Principle (Rosenthal, 1997).

Higher-order theories of consciousness vary depending on the nature of the higher-order process (percept-like, thought-like, indexical). Inner-sense theory proposes the existence of an inner-sense organ that produces higher-order percepts that would correspond to the content of conscious experience (Carruthers & Gennaro, 2020). Therefore, in the case of inner-sense theory, the first-order mental state is targeted by higher-order analog/non-conceptual intentional states. In contrast, the Actualist Higher-Order Theory (Rosenthal, 1997) states that the higher-order mental states giving rise to conscious experience are thoughts. So, conscious experience consists of first-order representations that are the object of higher-order thoughts.

Higher-order theories have not received as much attention as other theories of consciousness (Brown et al., 2019). However, there are recent theoretical proposals that are increasing the impact of higher-order theories. Lau (2019, 2022) suggested that the higher-order processes responsible for conscious experience are engaged in a process of Perceptual Reality Monitoring. This is an idea inspired by the concept of reality monitoring in memory (Johnson & Raye, 1981), the capacity to decide whether information had an internal source or an external perceptual source. The Perceptual Reality Monitoring hypothesis postulates a higher-order mechanism to determine whether a first order representation constitutes a reliable representation of external reality. First-order perceptual representations can be internally or externally generated. Furthermore, they may be unreliable. Thus, there is a need for a mechanism to determine when the first-order perceptual systems are engaged in reliable and externally originated activity. This mechanism is implicit and automatic. That is, it does not rely on a process of deliberate introspection. Finally, it is assumed that there is an overlap between the perceptual reality monitoring mechanism and the processes responsible for metacognition, the capacity to monitor our own mental activity. However, metacognition is assumed to be explicit and deliberate. Interestingly, for our current purposes, this hypothesis can be interpreted as postulating an FCC: perceptual reality monitoring.

Another interesting theoretical development was suggested by Fleming (2020) when investigating the mechanisms responsible for awareness reports. His central hypothesis is that "awareness is a higher-order state in a generative model of perceptual contents" Fleming (2020, p.2). This proposal assumes a predictive coding framework, in which the brain contains generative models of the world. These models usually have a hierarchical structure where top-down information consists of "hypotheses" about the state of the world. These "hypotheses" are updated by incorporating bottom-up information, based on the calculation of prediction errors (size of the discrepancy between the hypotheses and the incoming signals), from lower levels in the hierarchy. Fleming (2020) reported the results of several simulations showing an asymmetry between awareness reports about the presence or absence of stimuli. In presence reports, multiple potential states are nested in a common higher-order state. In contrast, in absence reports, the lack of stimulation results in fewer potential nested potential states. This asymmetry found in the simulations could account for some empirical findings observed when studying the neural correlates of consciousness. Specifically, it provides an alternative account of the increased global brain activity accompanying awareness reports.

Finally, it is worth mentioning one influential higher-order approach, the Radical Plasticity Hypothesis (Cleeremans, 2011). This hypothesis states that consciousness is the results of the brain learning about its own activity. It differentiates between the information *in* the system, important for sensitivity, and information *for* the system, relevant for awareness. Consciousness, in this hypothesis, would require meta-representations redescribing the information contained in first-order systems. Learning and plasticity are crucial processes for the generation of these meta-representations. The models presented in Cleeremans (2011) represent an exceptional exercise of how the philosophical ideas of the higher-order thought theory (Rosenthal, 1997) can be implemented and tested against empirical evidence.

Global Workspace Theory

One of the main features of the Global Workspace Theory is the contrast between two styles of processing. One is the vast, decentralized, and parallel information processing in the many specialized information processing systems of the brain. The other corresponds to a limited capacity system that stores the most immediate memories subject to voluntary control. This limited capacity system would allow the brain to have a "fleeting integrative capacity that enables access between functions otherwise separated" (Baars, 2017, p. 235). This workspace would play a crucial role in the functions of coordination, control, and integration of information.

The Global Workspace Theory fits within a broader hypothesis of consciousness, the global access hypothesis, see Figure 1, where consciousness is associated with the capacity of some information to have an impact on multiple consuming systems. This notion implies that one of the functions of the workspace is to broadcast information to many specialized information processing systems. A similar idea was suggested by Dennett (1993) in his "Multiple Draft Model". In this view, different versions of the state of the world, "drafts", compete to influence decision making systems. In contrast to the Global Workspace Theory, in Dennett's model, there is an attempt to eliminate the need of a central system to distribute conscious information.

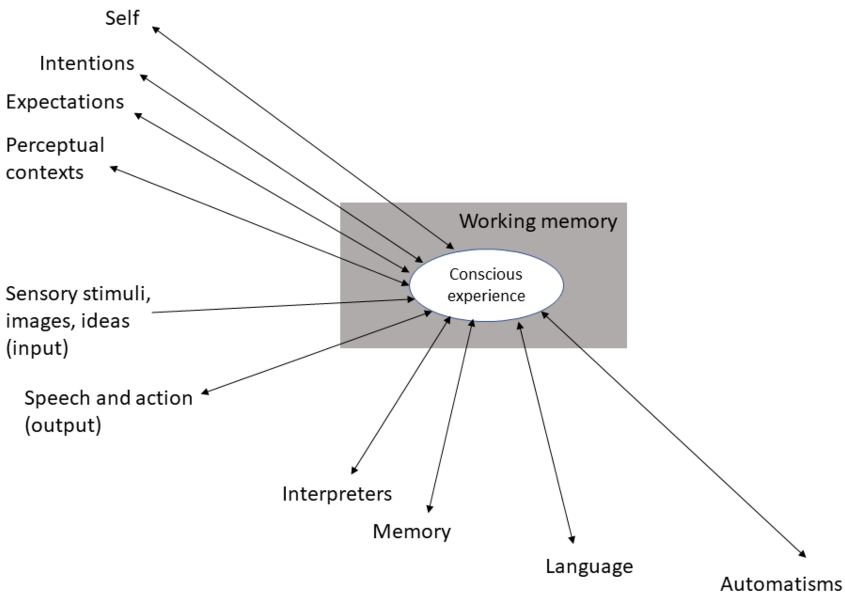

Figure 1. A schematic diagram of global access (adapted from Baars, 2017).

The workspace in the Global Workspace Theory could be understood as corresponding to the active components of working memory. In Baddeley's working memory model, the episodic buffer is supposed to be under the control of the central executive. The active representations targeted by the attentional mechanisms of the central executive could be understood as being part of the global workspace (Baars, 2003). As described above, the episodic buffer also played an integrative role by incorporating information from different modalities and from long-term memory. Therefore, the global workspace is crucial for one of the key higher-order cognitive functions: integration.

The Global Workspace Theory has inspired the creation of neural network models (Dehaene et al., 1998, 2003) accounting for important empirical findings, such as, the Attentional Blink (Raymond et al., 1992) (although, as discussed in Bowman et al (2022), it does not model very short lags well). In these models, the workspace acts as a bottleneck where selected information is amplified and gains access to specialize information processing systems. These models have inspired the Global Neuronal Workspace Theory (GNW) (Dehaene, 2014, Dehaene and Changeux, 2011, Dehaene et al., 2006, Dehaene and Naccache, 2001) suggesting that consciousness arises from the widespread dissemination of information across the cortex. This dissemination is facilitated by a network of interconnected high-level cortical regions, including the dorsolateral prefrontal cortex and the inferior parietal cortex, which together form a "global neuronal workspace." In this model, consciousness is believed to emerge when information becomes globally accessible and can be processed and integrated by multiple brain regions.

It would be an error to view the workspace as a higher-order process as conceived by the Higher-order Theories of Consciousness. Although, in both theoretical frameworks, first-order representations are not sufficient for consciousness, the main function of the workspace is not to produce an “inner awareness” (Brown et al., 2019). Rather, the global workspace is engaged in boosting and stabilizing sensory signals. What is broadcasted is the original signal and there are no new representations created.

Recurrent Processing Theory

Prior to describing the Recurrent Processing Theory (Lamme, 2003, 2010), it might be convenient to introduce a theoretical distinction about the nature of consciousness: the distinction between Phenomenal Consciousness and Access Consciousness. Up to this point in this Introduction, we have used the term consciousness to refer to information that can be declared or reported. One of the implications of our definition is the idea that consciousness is sparse. That is, the information available for report represents a tiny fraction of the information present in the system. Decades of studies on the capacity of short-term memory and WM have established some consensus about this view (Cowan, 2016). One of the consequences of identifying conscious representations with WM representations is that conscious access is restricted to the few elements contained in WM.

When presented with a brief visual display containing multiple elements, participants typically can only report a subset of those items (Sperling, 1960). However, it has been commonly argued that participants are aware that the display contained more elements than those they were able to report. It could be argued that their conscious experience of the display was not restricted to the reported items. Perhaps their conscious experience was richer, encompassing some/all of the items that were not reported. Therefore, to identify conscious content with verbal reports might be premature. Perhaps, a more apt account of this phenomenon is to differentiate between two types of consciousness. Access Consciousness would correspond to contents that can be reported and can have an impact on decision making processes, while Phenomenal Consciousness would identify the subjective experience of the access-conscious representations plus some phenomenal content that cannot be reported. Access Consciousness is sparse, since it only includes representations that passed the attentional bottleneck. In addition, it is often assumed that the nature of phenomenal conscious contents is analog/non-conceptual, while access-conscious contents could be conceptual. Block (1995) argued that this distinction between Access and Phenomenal Consciousness could have important theoretical implications (but see Cohen & Dennett, 2011).

One of the theories that has productively exploited this distinction is the Recurrent Processing Theory (Lamme, 2003, 2010). According to this view, a full account of consciousness cannot be restricted to the contents selected by attention and stored in WM. A theory of consciousness must account for Phenomenal Consciousness too. It follows from the above that there might be first-order representations giving rise to consciousness that, escaping the focus of attentional selection, are somehow different from the first-order representation that do not result in conscious contents. Then, what is the difference between conscious and non-conscious representations? Recurrent processing. It is stated that, for example in the visual modality, some functions (feature extraction, categorization) can be carried out during the first stage of visual processing, characterized by feed-forward (bottom-up) activation. In contrast, functions such as binding or perceptual integration only occur when there is recurrent processing (top-down and bottom-up simultaneous activity). In other words, some functions can recruit only local neural circuits, such as the primary visual cortex (Lamme, Zipser and Spekreijse, 1998), while other functions recruit global circuits that, through recurrent processing, result in conscious processing (Lamme, 2018). Phenomenal Consciousness would be the result of this second stage of processing, while Access Consciousness would correspond to those phenomenal conscious contents targeted by attention and stored in short-term memory.

Integrated information Theory

As in the Recurrent processing Theory, integration is postulated as one of the defining features of consciousness. However, the case for integration in IIT (Tononi, 2017b, 2017a) is not made on the basis of empirical findings or some theoretical solutions to the workings of our cognitive system. The case for integration originates in experience itself. One of the features that separates IIT from all the other theories of consciousness is that it starts from phenomenology. A series of self-evident (but see Bayne, 2018) axioms about experience sits at the core of the theory. These axioms are:

Intrinsic existence: our experiences are real.

Composition: experience has internal structure.

Information: each experience has a specific form.

Integration: the unity of consciousness. Consciousness does not seem to be reducible to non-interdependent components.

Exclusion: consciousness is definite, it has definite borders.

From these axioms, the theory establishes a series of postulates about any Physical Substrate of Consciousness: Cause-effect power upon itself (intrinsic existence), Internal structure (composition), Cause-effect power that is specific (Information), Cause-effect power that is irreducible (Integration), Cause-effect power that is maximally irreducible (Exclusion).

Of all the postulates, perhaps the most interesting for our current purposes (the FCCs) is the one related with Integration. First, because integration is an ingredient of other theories of consciousness and second because its postulate has been the most successful in terms of producing empirical predictions. That is, the theory suggests that integration of information varies and can be measured in physical systems. This modulation of integrated information can be quantified by Φ which refers to the "extent the cause-effect structure specified by a system's mechanisms changes if the system is partitioned (subdivided) along its minimum partition … by 'noising' the connections across it" (Tononi, 2017a, p. 246). The possibility of quantifying integrated information makes the theory testable.

Some of its more interesting predictions are related to the amount of integrated information in the human brain at any given time. For example, IIT predicts that the loss and recovery of consciousness should be related with the amount of integrated information. When consciousness is lost, low values of Φ are expected and when consciousness is regained a higher value of Φ is expected. In addition, IIT makes interesting predictions about the parts of the brain that contribute to consciousness. The connectivity within and between regions of the cortex makes it a splendid candidate to display high values of integrated information. In contrast, the structural properties of the cerebellum where the outputs of the different processing units are relatively independent from each other, make it a poor candidate to show high values of Φ and therefore unlikely to contribute much to our subjective experience.

4. Rapid Serial Visual Presentation (RSVP)

RSVP (Forster, 1970) is one of the most used experimental paradigms to study temporal attention and visual memory capacity limits. Typically, in RSVP, a sequence of stimuli is presented at the same spatial location at high rates (normally more that 8-10 items per second). Three main features characterize RSVP. First, the presentation time of each stimulus on the stream is very short. It is typically shorter than the duration of a single fixation. For example, when reading words (the stimuli used for the experiments described below), the mean fixation duration is 200-225ms. Second, each item has a forward mask (except the first) and a backward mask (except the last). Third, the sequences typically are formed by multiple items (rarely 6 or less) so the number of items is above the memory capacity limits of WM (Cowan, 2016). This paradigm has produced some of the most important effects for our understanding of the relationship between attention, visual short-term memory, and consciousness

The Attentional Blink (AB)

The AB (Broadbent & Broadbent, 1987; Raymond et al., 1992) is the most studied paradigm in the context of RSVP and temporal attention research. The AB, see Figure 2, is typically obtained when two task-relevant items (to-be-reported items) are inserted in close succession in an RSVP stream among distractors (task-irrelevant items). If the second target (T2) is presented between 200-600 ms after the first target (T1), the accuracy for reporting T2 (in fact, T2 conditional on T1 correct) is impaired relative to the accuracy for T2 when it is presented at later positions in the stream (more than 600 ms after the T1). In the search for the mechanisms responsible for the AB, a series of experimental manipulations have revealed some interesting findings:

Lag 1 sparing: When the T2 is presented immediately after (lag-1) the T1, the accuracy for reporting the T2 is at baseline levels (Potter et al., 1998). Interestingly, this preserved reportability of T2 comes with a cost for the accuracy for T1. In addition, the information about the order of the two events is degraded, and swaps are frequent at lag-1.

Spreading the sparing: When three or more targets are presented consecutively, accuracy for targets after the T1 is high even at lags 2 or 3 (Kawahara et al., 2006). A similar pattern is observed when comparing whole report (participants were asked to report all items in sequence) with partial report (with the classic AB manipulation). Accuracy is higher in whole report than in partial report (Nieuwenstein & Potter, 2006).

T2 Breakthrough effects: Accuracy performance for high salient T2s is relatively unimpaired. When personal names (Shapiro et al., 1997) or emotional stimuli (Anderson & Phelps, 2001) were used as T2s, performance is better relative to the blink observed with less salient T2s.

Priming: Shapiro et al., (1997) presented sequences of three targets embedded among distractors in RSVP. The key manipulation was that T2 and T3 could share the same identity (same letter in upper and lower case) or not. The results showed that missed T2s that shared the identity with T3 improved the accuracy for the T3. In the same article, they reported a semantic priming effect in the AB. They showed that a missed T2 word can facilitate the processing of a related word presented at the end of the stream.

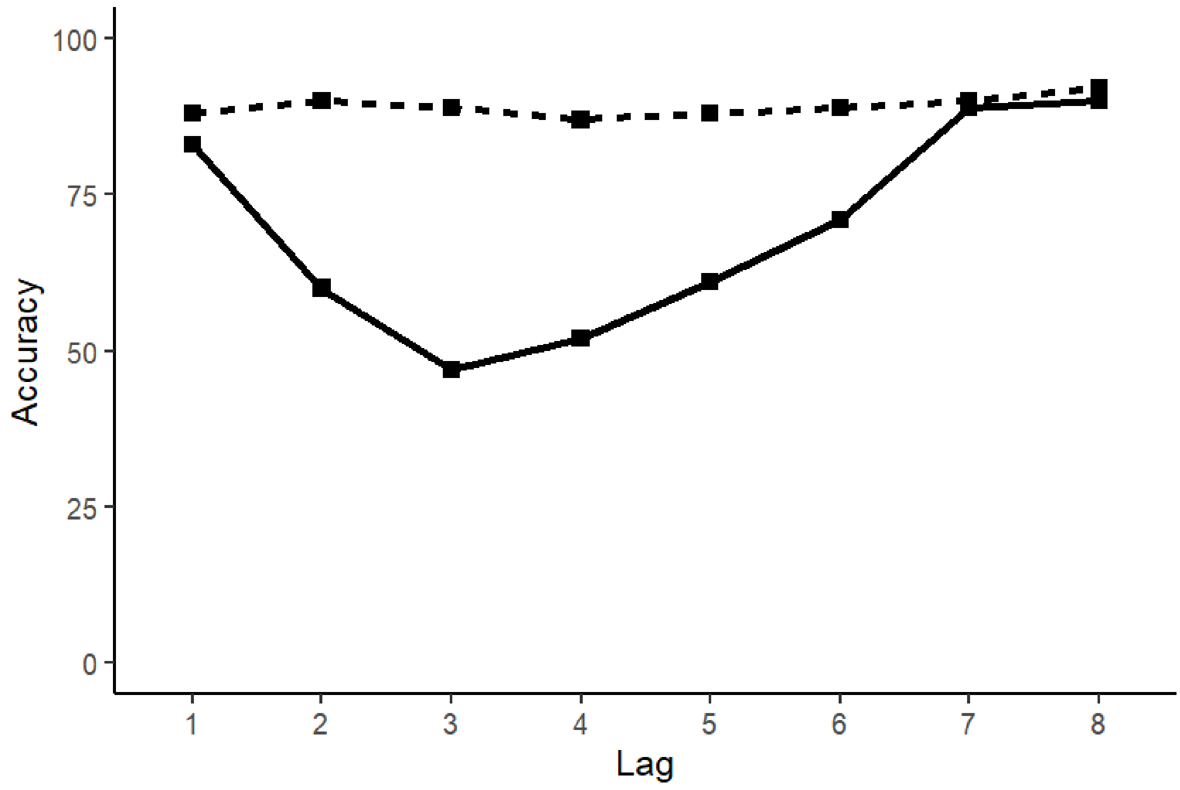

Figure 2. The basic "Attentional Blink" effect for letter stimuli (adapted from Raymond et al., 1992). Dashed line represents the T2 with no prior T1, and the solid line the T2 contingent on T1 report.

The AB and consciousness

The AB literature is particularly relevant for one of the topics discussed above: the relationship between attention and consciousness. Specifically, the AB evidence seem to support the notion that attention is necessary for consciousness (Cohen & Chun, 2017). Assuming that missed T2s in the AB implies a failure to consciously perceive the T2, then it might be concluded that without the adequate attentional focus, items in RSVP might fail to be consciously perceived. Interestingly the AB has been used to explore the relationship between subjective experience and reportability. Pincham et al., (2016) used RSVP and the AB paradigm to explore the possibility of working memory encoding without phenomenal consciousness. Participants rated the visibility of T2s presented at lags 1, 2, 3, 4, 5, 6 or 8. The results showed that accuracy followed the classic AB pattern with lag-1 sparing. Interestingly, the visibility results showed a very different pattern, with lower ratings for lags 1, 2 and 3 relative to lag-8. Thus, there is a clear dissociation between reportability and visibility at lag-1, with the blink curve for the latter (visibility) called the *experiential blink*. They argue that this supports the idea that we can "only perceive one thing at a time" since, although both the T1 and T2 are encoded simultaneously, there is not such simultaneity in the subjective experience of the participants.

Another contribution of the AB literature to the scientific study of consciousness comes from priming effects (Shapiro et al., 1997) and the breakthrough of the T2 (Anderson & Phelps, 2001). These effects show that items in the AB window undergoing deep processing and therefore the attentional bottleneck is located at a late stage. That is, items in the AB window are processed deep enough to activate lexical and semantic representations (Shapiro et al., 1997) that might facilitate the processing of subsequent items. Therefore it can be argued that the filtering of the attentional bottleneck acts at later stages of processing. These results support the idea that basic functions of visual processing, such as feature extraction, categorization, and activation of the semantic properties of linguistic stimuli can be performed on the basis of nonconscious representations.

Repetition blindness (RB)

RB (Kanwisher, 1987) refers to the impaired reportability of a T2 when two targets that share the same identity are presented in close succession in RSVP. This result contrasts with the expected facilitation from repetitions observed in priming and memory studies. Kanwisher (1987) showed that participants failed to detect repeated words (Experiment 1); and also failed to report the second instance of a repeated word in a sentence (Experiment 2). Interestingly, a facilitation effect on the second instance of a word was found when participants were asked to report exclusively the last item of the stream. Performance was better in streams with repeated words relative to the unrepeated condition.

Chun (1997) summarized the similarities between the AB and RB, and conducted a series of experiments exploring the differences between the two effects. The AB and RB seem to reflect capacity limitations in attentional processing. Both effects require attention to the T1, and are dependent on the distance between the T1 and T2. That is, at lag-8 it is unlikely to find an impairment in the reportability of the T2 in any of the two paradigms. Finally, the AB and RB require a rapid presentation of items (8 or more per sec). The results of the experiments in Chun (1997) showed that increasing the discriminability of targets in the sequence (letters presented among keyboard symbols) attenuated the AB but not the RB. In contrast, increasing the discriminability between the two targets (using different colours for T1 and T2) resulted in a reduction of RB and a normal AB. They suggested that a Type/Token Framework (see below) could account for the similarities and the differences between the AB and RB.

Detection, Identification and Recognition in RSVP

> "I took a speed-reading course and read War and Peace in twenty minutes.
> It involves Russia."
> — Woody Allen

Perhaps one of the most striking and yet consistent results obtained with RSVP is the excellent performance of participants at detecting salient stimuli. In Potter (1976), participants were asked to search for a target picture in RSVP sequences with stimuli presented at different presentation times (113, 167, 250, 333 ms). In the detection condition, prior to the presentation of the RSVP sequences, the participants were presented with a target picture or a brief description (e.g., “a road with cars”). A third group of participants (the recognition group) were asked to simply view the sequences of pictures, and were tested with a recognition question presented at the end of each stream. The results showed an excellent performance for detection, even at the fastest presentation rate (accuracy of 64% for the name target, and 73% for the picture target). In contrast, performance for recognition at the fastest rate was an 11% accuracy (corrected for guessing). This high performance in a detection or identification task in RSVP has been observed with letters and digits (Chun and Potter, 1995), pictures (Intraub, 1981), and words (Potter et al., 2002). Interestingly, above chance Detection performance has been observed even with pictures presented at 13 ms (Potter et al., 2014).

One of the potential explanations for the high detection performance in RSVP could be that expected/searched-for stimuli may facilitate target identification by lowering the targets’ recognition threshold. Potter (1976) tested that hypothesis by briefly presenting a single masked picture and testing participants in a recognition test. They found a high recognition performance (80% accuracy) for a presentation time of 120 ms. Since the target was not searched for and recognition performance was equivalent to that obtained for detection in RSVP, they concluded that the “expectancy hypothesis” was not responsible for the high detection performance observed in RSVP. Instead, they argued that most of the items in RSVP sequences are deeply processed up to the conceptual/semantic level. They proposed that a processing and memory system, called Conceptual Short-Term Memory (CSTM), was responsible for briefly storing conceptual information about the stimuli. CSTM would play a role in perceptual integration and in integrating active representations into larger structures (e.g., sentences). Interestingly, if CSTM representations fail to be incorporated into larger structures then they are rapidly forgotten (and perhaps not consciously perceived). Potter (2012, p.8) addressed the issue of the relationship between CSTM and consciousness, as follows: "Is CSTM conscious? The question is difficult to answer, because we have no clear independent criterion for consciousness other than availability for report. And, by hypothesis, report requires some form of consolidation; therefore, only what persists in a structured form will be reportable. Thus, while the evidence we have reviewed demonstrates that there is conceptual processing of material that is subsequently forgotten, it does not tell us whether we were briefly conscious of that material, or whether the activation and selection occurred unconsciously." The experiments reported in this thesis aimed to address exactly this issue.

Bowman et al. (2013) suggested that the concept of a Subliminal Salience Search was useful to describe the capacity to detect salient stimuli in RSVP. When presented with an RSVP stream, it can be said that our perceptual system is effectively searching for one or more items matching one of the templates determined by the task. The term salience refers to the properties of a stimulus that make it stand out from the rest. These properties could be intrinsic; that is, could be independent of the subject's task-set or goals. Or, instead, saliency could be determined by the instructions or the structure of the task. Finally, the search process is said to be subliminal in the sense that most of the items in a sequence are processed deeply enough to enable the saliency decision but still fail to be reported. As mentioned above, recognition performance is extremely low in RSVP. Consequently, there is a huge gap (that we call the Missing Memory Imprint, see Chapter 4) between detection and recognition performance. If detection accuracy is typically above 75%, then it can be said that, at least, 75% of items are processed to the extent of enabling the saliency decision. In contrast, reportability is nowhere near that level. We will argue in Chapter 4, in line with the Subliminal Salience Search idea, that the Missing Memory Imprint is the consequence of two different stages of processing, and that detection is executed on the basis of nonconscious representations.

It has been argued that a low recognition performance for items presented in RSVP does not imply that there are no memory traces left by those items. In line with this argument, Endress & Potter, (2014) showed that simply by repeating a critical item in RSVP there is an increase in recognition performance. They presented the same picture embedded among distractors in 1, 2, 4, 8 or 16 RSVP trials and tested memory for targets with a recognition task. They showed an increase in performance with more repetitions. They concluded that memories are consolidated through repetition. In agreement with this idea, Thunell and Thorpe (2019) showed that not only did recognition performance increase with repetition but also participants have more chances of detecting a repetition in RSVP as the number of repetitions increased.

In contrast, Bowman et al. (2014) showed that repeating an item 50 times in different RSVP trials in one single session did not result in participants noticing the repetition. In this experiment, participants were asked to search for a target name. In addition, participants were asked to find a name that was going to be presented often during the experiment. The results showed that they failed to find the repeated name. It is reasonable to assume that if every stimulus in the RSVP streams had been consciously perceived, participants would have been considerably more successful at finding the repetitions (for example, at 5 words per second participants detect all repetitions, Kanwisher, 1987). One tentative conclusion -and the inspiration for the current work- is that participants failed to detect the repetitions because most of the items in the RSVP streams were not consciously perceived. We revisit this idea in Chapters 2 and 6.

5. The Type/Token framework for visual processing.

Rose is a rose is a rose is a rose

Gertrude Stein

Consider the verse above (Wetzel, 2018). How many words are in this verse? There are two possible answers to this question. One might say that there are three different words ("is", "a", "rose"). But also, it is equally correct to say that there are ten words. One solution to this ambiguity is to say that there are three "types" and ten "tokens". Types are abstract, while tokens have a spatio-temporal location. Our cognitive system manages type and token information quite effectively. Questions tapping Type information (e.g., Does the word "rose" appear in the verse?) can be easily answered. The same goes for Token information questions (e.g., How many times the word "rose" appears in the verse). However, arguably, dealing with Type or Token information seems to require very different processes considering that Types consists in abstract and "timeless" information while Tokens demand the processing of spatio-temporal information. In fact, this distinction between Type and Token information has proven useful to account for empirical findings in the field of temporal attention.

The token individuation hypothesis (Kanwisher, 1989)

Kanwisher (1989, p. 119) argued that "repetition blindness occurs when words are recognized as types but not individuated as separate tokens", the Token Individuation Hypothesis. As described above, Repetition Blindness (RB) refers to the impairment at reporting the second instance of a repeated stimulus in RSVP. For example, when presented with the sentence "Her jacket was *red* because *red* is conspicuous" in RSVP, participants frequently omitted the second instance of the word "red" in a free-recall test. Kanwisher (1989) presented a model of visual information based on the Type/Token distinction. The model differentiates between the categorical information "Types" of words and the episodic information "Tokens" about the serial order of events. Type and token information would be encoded concurrently but in different domains. Each stimulus in RSVP would activate the corresponding type and token nodes and a link would connect the activated type node with the corresponding token. They called this process of token assignment "token individuation".

RB would occur because once a type has been token individuated, it would not be available for a second token individuation. So, when the second instance of a stimulus is presented, its type would be activated but the link to the corresponding token would be inhibited. This implies that there is an absence of temporal information for the second instance of the stimulus, and accordingly the "subject will only be aware of one occurrence of the repeated word" (Kanwisher, 1989, p.134).

A two-stage model of the Attentional Blink (Chun & Potter, 1995)

In their account of the AB effect, Chun & Potter (1995) proposed a two-stage model that, although they did not use the terms "type" or "token", share some features with the Token Individuation Hypothesis. They suggested that two-stages of processing were needed to account for the AB. The first stage would consist in the feature extraction and categorization of the stimuli, akin to Potter's Conceptual Short Term Memory. They suggested that in RSVP, the categorical identity of most of the items is briefly available (Chun & Potter, 1995). This is not far from the notion of Types introduced by the Token Individuation Hypothesis. In addition, they proposed a capacity-limited second stage of processing. This stage is needed because the representations of the first stage are fragile and cannot be the basis for report. The second stage would carry out the function of memory consolidation. Importantly, their model reproduces the AB because when a T1 is selected and is being consolidated in stage 2, no subsequent items would be processed beyond stage 1.

A Type/Token Framework for Visual Processing (Chun, 1997)

As described above there is a Type/Token account of RB (Token Individuation Hypothesis) and a 2-stage account of AB (Chun & Potter, 1997). Chun (1997) hypothesized that there was sufficient commonalities between the two ideas to suggest a common framework that could account for both effects. A 2-stage model was proposed with a first stage primarily responsible for the processing of type information and a second stage in charge of producing object tokens containing episodic information. Object tokens are a similar idea to the construct of Object Files (Treisman, 1988), internal representations that integrate long-term categorical information with spatio-temporal information. Object tokens are supposed to sustain conscious reports and guide action. Interestingly, Chun (1997) suggested that while object tokens would be the output of the system, the model requires a third type of units: Spatiotemporal Tokens that would encode pure episodic information (based on the discontinuities between each stimulus in RSVP). So, an object token requires a successful association between a type and a spatiotemporal token, a process called tokenization. In this model, AB and RB represent failures within the stage of tokenization. As in the 2-stage model of Chun & Potter (1995), the model "blinks" when the second stage is engaged in the consolidation of a selected representation from the first stage. In this case, the second stage is the tokenization process. RB is caused by a failure to tokenize two types, as in the Token Individuation Hypothesis.

The simultaneous type/ serial token model (Bowman & Wyble, 2007)

The Type/Token Framework served as inspiration for one of the most successful neural network models accounting for the AB, the Simultaneous Type\ Serial Token (ST^2) Model (Bowman & Wyble, 2007), see Figure 3. A key aspect of this model is that WM is not just about the consolidation and maintenance of representations, but it is also engaged in the assignment of episodic information ("temporal contexts"). According to this view, one of the main functions of WM is to encode the order of events. In addition, WM processes must be capable of dealing with the repetition problem, that is, in order to detect a repetition, the system must have the capacity of associating the same representation with two different episodic contexts. One of the core ideas of the model is that there is a "cost" in assigning these episodic contexts, and the AB is a manifestation of that cost. Importantly, the model makes use of the differentiation between types and tokens to implement these functions.

ST^2 is a 2-stage model of temporal attention (Chun & Potter, 1995). Stage 1 implements visual processing functions up to semantic categorization. It is at this stage that types are extracted. Types refer to the semantic, categorical, and visual features of stimuli. There are different layers of processing units in stage 1. The last (most distant to the Input layer) of these layers is the Task Filtered Layer. It acts as a salience filter enhancing relevant items and discarding irrelevant items. Importantly, one of the main features of this stage is parallel processing, that is, multiple items can be processed simultaneously in stage 1.

Stage 2 is the process of WM consolidation. This process is sequential as the system associates types with episodic contexts. These episodic contexts are implemented by units called tokens, and they encapsulate temporal order information. Thus, the main task of

WM encoding is the tokenization process, the association of a particular type with the corresponding token. A key feature of the model is that once a token has been allocated, type information is no longer needed for WM maintenance. This implies that the same type can be assigned to a later token. Thus, the system is capable of detecting repetitions, since the same categorical and perceptual information may end up associated with two or more different episodic contexts.

The transition from stage 1 to stage 2 is mediated by an attentional component. Once an item has been identified as task-relevant, a Transitory Attentional Enhancement (TAE) mechanism elevates the activation level of the processing units of the relevant item. This mechanism enables the tokenization process for that item and remains inactive until the tokenization process is over. Consequently, the TAE ensures that a token is associated to one single type (with the exception of a T2 at lag-1). If a second task-relevant item is presented while the tokenization of the T1 is ongoing (and the TAE is shut off) then the T2 type will fail to reach stage 2. This failure in the WM encoding of a T2 is responsible for the low reportability of T2s in the AB. In contrast, if the T2 arrives after the tokenization of the T1, it can be tokenized, and it will thus become available for report. In virtue of this mechanism, it can be said that the AB is the cost that the system pays for imposing a temporal discretization on the input.

It is worth mentioning here a final component of the model, the Binding Pool. This is a set of nodes that encode the associations between Tokens and Types. Every combination of token and type results in a unique pattern of activity in the Binding Pool. This feature enables the token units to act as pointers, since the information about the type and token association is encoded in the binding pool.

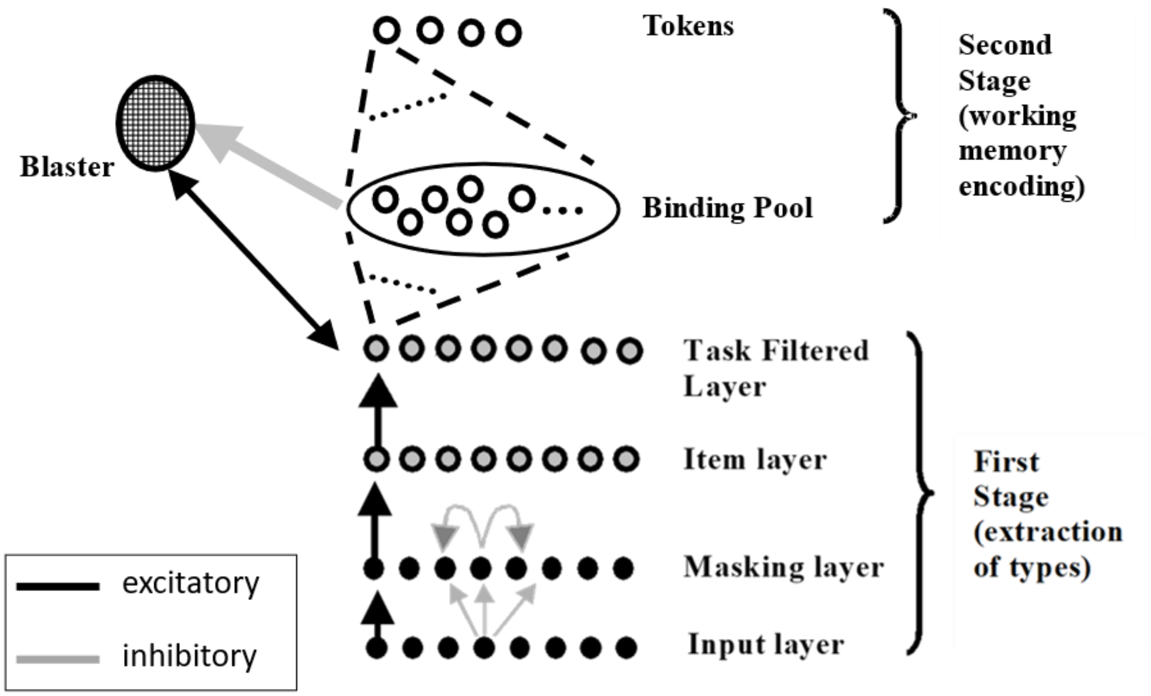

Figure 3. Simplified depiction of the ST^2 model. In the initial stage of processing, units operate concurrently to handle both featural and categorical information. Subsequently, units in the second stage engage in a sequential process to encode information into working memory (WM). The transition from the first stage to the second stage is facilitated by a Transitory Attention Enhancement mechanism, which is realized through a specific unit referred to as the "blaster." This blaster unit plays a pivotal role in enhancing attention and enabling the efficient encoding of salient items into WM.

The model successfully accounts for many of the effects associated with the AB. A review of specifics mechanisms that allow the model to replicate some of the most relevant empirical findings is beyond the scope of this Introduction. However, perhaps it is worth to describe how the model accounts for the lag-1 sparing effect. As mentioned above, lag-1 sparing is the preserved reportability of a T2 presented immediately after the T1. Lag-1 sparing, in the model, is the consequence of a T2 participating in the Transitory Attentional Enhancement mechanism activated by the presentation of a T1. Because of this, the T2 is tokenized alongside the T1, becoming associated to the same token. The implication of this is that order information is not properly encoded to differentiate T1 from T2. As a result, swap errors are to be expected for these items.

6. The tokenized percept hypothesis.

One recurrent theme in the different versions of the Type/Token framework is that conscious reports are possible only when representations become associated with an episodic context. Without this temporal contextualization, stimulus representations are irremediably fragile, and they are quickly forgotten. Being anchored to an episodic context is how representations gain durability and episodic distinctiveness. The duration of representations is crucial for conscious reports. Fleeting representations are inevitably not available for processes, such as report, that normally are engaged after the extinguishing of physical stimulation.

Kanwisher (2001) explicitly promoted the Type/Token framework to a hypothesis about conscious perception. She argued that “awareness of a particular perceptual attribute requires not only the activation of a representation of that attribute, but also individuation of that perceptual information as a distinct event” (Kanwisher, 2001, p. 107). A fully fledged conscious percept is a discrete object that appears in a concrete spatiotemporal location. According to this view, attention is necessary for consciousness (see above and Cohen and Chun, 2017). The same attentional resources to those recruited for binding (Treisman & Gelade, 1980) would act to associate perceptual attributes to episodic contexts.

The link between episodic encoding and conscious perception is not exclusive of the Type/Token framework. Droege (2009) made a distinction between the mental states that represent their contents *as now*, and those that simply *occur now*. “Events occurring now need not be represented as now” Droege (2009, p. 79). Conscious states necessarily represent a concrete point in time, now. In order to do that, temporal relations need to be encoded. Similarly, in their theory of Explicit and Implicit knowledge, Dienes & Perner (1999) also emphasized that factuality of an event needs to be explicitly represented to produce a conscious percept. For example, under subliminal conditions, some properties of the stimulus are explicitly represented (its categorial information), but its factuality (that there is a particular stimulus event) remains implicit.

We think that these ideas are largely consistent, and they converge on what we call the *Tokenized Percept Hypothesis*. It states that one of the functional correlates of consciousness is the process of episodically tagging experiences. According to this view, the subconscious is representationally rich, but “episodically blind”. There is no explicit knowledge of temporal order that is nonconscious. In terms of Droege’s view, nonconscious representations occur *now* but they are not encoded *as now*. The subconscious is the realm of the “what” but not the “when”. Features can be extracted, categories activated, and semantics triggered subconsciously. In contrast, consciousness recruits episodically rich but content-independent information, tokens. A conscious percept is, in our proposal, a tokenized percept; that is, a perceptual and/or categorical representation associated with a particular episodic context.

According to this hypothesis, nonconscious processing is parallel. The nonconscious type stage of processing is characterized by the extraction of perceptual features and the activation of categorical information. Information processing at this stage is parallel:

- multiple items can be processed at the same time (Bowman & Wyble, 2007).

- "vertical" parallel processing: Earlier and later layers of processing interact (McClelland & Rumelhart, 1981)
- "horizontal" parallel processing: Active representations have an impact on the activation level of related representations (Collins & Loftus, 1975).

In RSVP, it is assumed that almost all items activate the correspond type information. As Potter (1976) showed, the first 100-120 ms of processing might be enough to activate conceptual representations. This explains the excellent performance in detection and identification tasks. We add that type information that is not intrinsically salient or task-relevant would remain as nonconscious content. Consequently, the T2 in the AB and many of the distractors would not be consciously perceived.

The token stage of processing is serial. The seriality of this stage processing determines its function: to impose episodic discreteness on the input. This is best illustrated by our capacity to detect repetitions. To detect a repetition, a second instance of a stimulus needs to be processed in the context of a previously encoded first instance. The first instance is part of the episodic context for the second instance. The seriality and content-independence of tokens enable this process.

The Tokenized Percept Hypothesis shares some features with the Higher-order theories of consciousness. First, the first-order representations, types, are not sufficient for consciousness. Second, tokens are higher-order units, not only because they target first-order representations but also because they carry information about the state of the system. If tokens indicate temporal order, they need to be sensitive to past states of the system. In ST^2, availability of token2 not only means that there is a token ready for binding but also that token1 has been allocated. Consequently, the Tokenized Percept Hypothesis states that memory is necessary for consciousness. If a token being allocated carries information about the past, memory cannot be disentangled from conscious perception.

7. Overview of the thesis

That many of the items presented in RSVP cannot be reported is a well-established finding since Potter (1976). Many of the empirical findings and theoretical ideas described above are consistent with an interpretation in which most of the items in an RSVP stream are not episodically encoded. However, there is a recurrent objection to this interpretation that argues that failures of reports do not imply necessarily a failure of episodic encoding. Indeed, most of the items in RSVP could be reportable (with the implication of being consciously perceived) for a brief period of time but they are quickly forgotten (see Rees et al., 1999, for evidence against this in an Inattentional Blindness experiment). Items could fail to be reported simply because too much time passed between physical stimulation and the memory test.

Arguably, evidence showing that many RSVP items result in episodic representation admits two different interpretations: a) most of RSVP items are consciously perceived; b) most of RSVP items are not consciously perceived and yet they are episodically encoded. Both interpretations are difficult to conciliate with the Tokenized Percept Hypothesis. The "everything conscious" interpretation challenges the notion of two stages of processing favouring a graded view of episodic encoding. On the other hand, the "unconscious and yet episodic" attacks the core claim of the hypothesis, that episodic encoding is one of the functional correlates of consciousness.

In the light of the above, it seems essential to explore to what extent RSVP items result in episodic representations. The experiments reported here were designed to do that. Chapter 4 addresses one of the shortcomings of a “standard” memory test in RSVP: the possibility of forgetting. We tested the memory for the last item of an RSVP sequence. By reducing the retention interval between presentation and test we aimed to provide a direct test of the reportability of RSVP items.

The rationale for the other Chapters is to explore episodic memory effects in RSVP. In Chapters 2 and 6 we explore the Repetition effect (Ebbinghaus, 1964). In Chapter 3, we tested Recency effects in Recognition and Free-recall. Finally in Chapter 5, we tested Proactive Interference effects. These effects observed in RSVP at fast rates of presentation were compared with performance obtained at much slower rates (where it is reasonable to assume that every single item is consciously perceived). If the effects observed in RSVP at fast rates are comparable to those reported in the episodic memory literature, the hypothesis that many of the items in RSVP are episodically encoded would gain credibility. In contrast, if the observed effects are compatible with the idea that many items fail to be episodically registered, the Tokenized Percept Hypothesis will have passed its first test.

Chapter 2: On the limits of evidence accumulation of the preconscious percept

Avilés, A., Bowman, H., & Wyble, B. (2020). On the limits of evidence accumulation of the preconscious percept. *Cognition*, 195, 104080. https://doi.org/10.1016/j.cognition.2019.104080

1. Introduction

It is uncontroversial to believe that memories for a stimulus are strengthened through repeated experience of that stimulus. Indeed, Endress and Potter (2014) have explicitly shown this to be the case, even for briefly fixated objects. However, a key question is the role of conscious perception in this process. In this respect, the scenario presented in Case 1 in figure 1 is not controversial; that is, if a stimulus is consciously perceived, some sort of "trace" (whether activation-based or synaptic) of that stimulus would typically form after the moment of awareness. Furthermore, such a trace could accumulate over repeated presentations, i.e. there would be *evidence accumulation*.

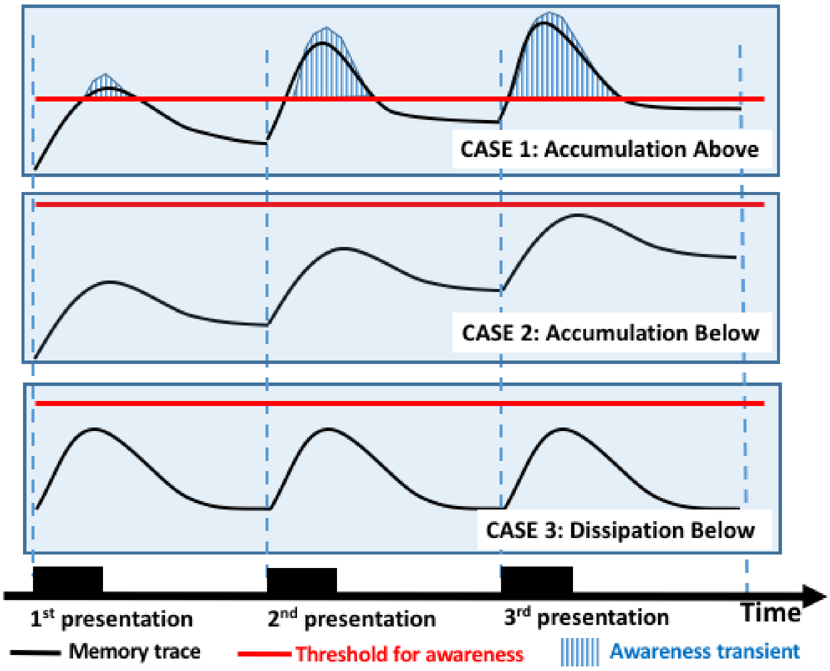

Figure 1: three theories of how the brain responds to repeated presentations. The stimulus sequence is shown in black as three presentations of the same stimulus. The awareness transient reflects the conscious experience of the presented stimulus. Three accumulation regimes are shown. CASE 1: evidence accumulates across presentations, each of which yields an, if only brief, conscious percept. CASE 2: evidence accumulates without conscious percepts. CASE 3: evidence dissipates between presentations, none of which generate a conscious percept.

What happens though, in the case of stimuli that do not reach awareness is less clear. That is, would it be possible for the trace of an item to strengthen, i.e. for evidence to accumulate, through repeated presentation if none of those presentations induced a conscious percept, such that conscious perception becomes more likely with further presentations, as in Case 2 of Figure 1? In contrast, it could be the case, as shown in Case 3, that stimuli that do not reach a threshold of awareness do not leave a trace that accumulates over multiple exposures.

This is the question we consider in this paper, viz, we seek to determine whether the brain behaves as Case 2 or Case 3 of figure 1. In other words, do representation traces for stimuli registered below the the awareness threshold, dissipate back to baseline so quickly that there is effectively no evidence accumulation for the preconscious percept? (How these different modes of evidence accumulation relate to relevant phenomena in the literature, such as subliminal priming, is considered in the Discussion section.)

To answer this question, one needs a means to present a lot of stimuli in such a way that many do not cross the awareness threshold. The natural way to do this is with Rapid Serial Visual Presentation (RSVP), for which it is known that only a small subset of the presented stimuli are reportable, or indeed recognisable (Bowman et al., 2013; Bowman, Filetti, Alsufyani, Janssen, & Su, 2014; Potter, 1976).

In fact, there are previous studies that have considered the progressive strengthening of memories with RSVP (Albrecht & Vorberg, 2010; Endress & Potter, 2014; Subramaniam, Biederman, & Madigan, 2000). For example, Endress and Potter (2014) reported that images (and words) presented more often across a number of Rapid Serial Visual Presentation (RSVP) streams were recognized more accurately in a final recognition test. Given that, as previously discussed, in RSVP studies, participants often show very poor recognition performance, Endress and Potter's findings are open to the interpretation that some items leave memory traces that gain strength through repetition, despite not being consciously perceived. That is, although we will ultimately argue against this position, Endress & Potter's findings open the possibility that Case 2 of figure 1 obtains and memories accumulate through repetition for the pre-conscious percept.

However, some studies have reported that most stimuli presented in RSVP do not display consolidation/evidence accumulation through repeated exposure. Subramanian, Biederman and Madigan (2000) presented participants with RSVP streams of drawings of objects. Participants were instructed to search for a target image. Crucially, some of the non-target pictures were presented 15 times on average before becoming a target. The results revealed an absence of a repetition effect; that is, participants were not better at detecting the targets that had previously been repeated relative to those presented once. In a similar vein, Bowman et al., (2014) presented participants RSVP streams of first names and they were instructed to search for a Fake Name (a name they were pretending was their name). In addition, in Experiment 3, they were instructed to search for frequently presented names; that is, they had to search for names simply on the basis that they were repeated. Importantly, these repeating names were presented as often as the Fakes (up to 50 times). The behavioral results (recall and recognition test at the end of the experiment) indicated that participants found it very hard to identify the repeated names; and consistent with this, in ERP findings, the Fake generated a clear P3 that was absent for the repeating name. As in Subramanian et al. (2000), participants were actively searching for an additional item (Target in Subramanian et al, Fake in Bowman et al), which arguably could hinder the encoding process of repeating items.

Taken together, the above described findings are far from offering a coherent picture. On the one hand, some studies raise the possibility that evidence accumulation is possible for non-retrievable (and thus not consciously perceived) stimuli presented in RSVP[2]. This would suggest that items naturally elicit (graded-strength) memory traces. On the other

[2] Our focus in this paper is specifically on access awareness (Baars, 2002; Block, 2007) which means that conscious perception can be associated with retrievability; that is, report is taken as an indicator of awareness.

hand, other studies suggest that the capacity to form memory traces from the fleeting representations of items in RSVP is very limited. That is, despite the serial presentation of items in RSVP, unless an item is processed to the point that it reaches a state of awareness, evidence for it would not accumulate.

One can view the present study as shedding light on the nature of the discrepancies between the described findings. If the presentation of stimuli in RSVP results in memory representations of gradual strength, these representations may increase in strength with every repetition, facilitating retrieval. On the other hand, the encoding of items could follow a bottleneck behavior: a small number of stimuli would (enter consciousness and) be stored in stable WM representations, while the vast majority would not reach that stage. Importantly, the “missed” stimuli would not benefit from successive repetitions, since residual information would not survive the presentation of new items.

In the previous studies that have considered the progressive strengthening of memories with RSVP (Subramanian, Biederman & Madigan, 2000; Endress & Potter, 2014), a repeating stimulus occurred incidentally in RSVP streams through the course of an experiment. Then a recognition test on the stimulus was inserted after a certain number of repetitions. This previous work, though, was not specifically focussed on *preconscious* evidence accumulation, which is our interest. That is, they did not probe memory in such a way that they could identify *the first time* a stimulus was seen as repeating, leaving the possibility that their recognition reports could have arisen after at least some previous presentations were consciously perceived, i.e. Case 1 in our figure 1. As a result, we have had to employ a somewhat different experimental paradigm to these previous studies. In particular, we could not rely on recognition memory tests at the end of each of a number of RSVP streams, since such a test could reveal the identity of a repeating (target) stimulus whether it had or had not been perceived in a stream to that point, thereby consciously priming its future perception. This would confound any test of an *intrinsic* below threshold build-up of evidence with repeated presentation.

We are interested in isolating the *first* instance at which a (non-primed) repeating item is consciously perceived as repeating. To obtain such a test, we have run an RSVP repetition experiment, where we instruct participants that a repeating item will be presented, but we do not identify it and participants are required to search for it simply on the basis that it repeats.

With this approach, we can test what turns out to be the key property for us, which is that (first) detection of repetition is invariant to the number of prior presentations, with the following procedure.

1)We determine the first instance at which an item is seen as repeating.

2) We assess whether the probability of this first instance is invariant to the number of prior repetitions; by determining the conditional probability of seeing a repetition, given that it has not been seen as a repetition before.

If this conditional probability, which we call the *first seen as rep. probability*, is indeed invariant across repetitions, there is no evidence accumulation before a repetition is first seen. This can be illustrated with reference to figure 1, where the accumulation inherent to Case 2 would ensure that the probability of first seeing (and then seeing as a repetition) would increase with each sub-threshold registration of a stimulus. This is because the distance to threshold would be reducing on each registration. In contrast, in Case 3, the probability of first seeing (and then seeing as a repetition) would not change with (below threshold) registration of a repeating stimulus. This interpretation of Cases 1 and 2 is confirmed in simulations in the appendix.

More specifically, we have designed two different tasks using RSVP of streams of words: (1) a Repetition task and (2) a Detection task. In the first of these, the repetition task, participants were instructed to search for a word that was repeated across trials/ streams. They were informed that they were going to see several streams of words and they must search for the repeated word. At the end of each trial/ stream, they had to answer the question "Which one was the repeated word?". We varied the presentation time in two experiments (from an SOA of 17ms to 533ms). In this task, we tested to what extent participants can detect repeated items inserted in different RSVP trials. This enabled us to examine to what extent contents generated in RSVP can accumulate evidence through repetitions, with a key test being invariance to prior presentation of the first seen as rep. probability. To do this, we examined whether the number of presentations of a word increased the probability of first detecting the repeated word.

While, as we have discussed, the experimental paradigm we employ here is somewhat different to that used in (Subramanian, Biederman & Madigan, 2000; Endress & Potter, 2014), we believe our experiment and theirs are comparable. In particular, our experiment can be seen as a generous test of the evidence accumulation question, since we are *instructing participants to look for repetitions*. That is, if evidence does not accumulate preconsciously for repeating items, when participants are explicitly instructed to look for such items, it seems unlikely that is would incidentally, when participants are not instructed to look for them. Incidental build up is the approach in the Subramanian, et al and Endress & Potter experiments.

The results of the Repetition Task were compared with those of a Detection task, which effectively served as a baseline to compare against. In the Detection task, participants were instructed to search for pre-specified target words. This enabled us to explore to what extent participants can search for stimuli in RSVP on the basis of their perceptual properties. With the comparison of the two tasks, we aimed to illustrate the time-course of two fundamentally different types of search: one based on the perceptual features of task relevant items (Detection) and the other based purely on frequent occurrence (Repetition). Our findings will confirm that, consistent with previous work (Bowman et al, 2013; Potter, 1976;), the brain is exceptionally good at searching for pre-specified items, as per our detection task. In contrast, searching on the basis that an a priori unknown item occurs frequently, as per our repetition task, is much harder.

2. Method

2.1 Experiment 1

Participants

21 undergraduate students of the University of Birmingham took part in Experiment 1 in exchange for course credits. All were right handed, native English speakers and had normal or corrected-to-normal vision. The experiment conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Materials

36 English nouns were selected from the English Lexicon Project database (Balota et al., 2007) to serve as Targets. Additionally, 1800 English words were selected as Distractors from (Warriner, Kuperman, & Brysbaert, 2013). Two lists of Targets were created so that participants could perform the Detection and Repetition task on a different set of Targets. Both lists of Targets were controlled for word frequency, concreteness and

emotional valence. All targets and distractor words had 6 letters and none of them were proper nouns. The same set of Distractors was used in both lists.

Procedure

RSVP streams were presented on a 24’’ LCD screen (refresh rate: 60Hz, resolution: 1600 x 1200) using custom PsychToolbox scripts running under Matlab 2016a. Stimuli were 16-point white (75% white) characters on a dark background (25% white). Participants were seated at 60cm from the screen.

The experiment was conducted individually in a quiet room. Each experimental session was divided into two tasks (Detection Task and Repetition Task). The order of tasks was counterbalanced (15 participants performed the Detection Task first and 16 participants performed the Repetition Task first). Each task consisted of 18 blocks (3 blocks per SOA condition: 33, 133, 233, 333, 433 and 533ms) presented in random order. The target was the same within a block and different across blocks.

Detection Task

Each block consisted of 10 trials. At the beginning of the block, participants saw the instruction “Search for the word:” in the upper part of the screen along with the target word presented at the centre of the screen for 2 seconds. Each trial started with the centred presentation of the fixation cross (+) for 500ms. Then a stream of 10 words was presented and a final item “######” at the end of the stream. The trial ended with the question “Have you seen the target word?”. Participants responded to the question by pressing predetermined “yes” and “no” buttons. The target word was presented in 5 of the

10 trials of each block and the remaining 5 trials were target absent trials. The target word could be presented in any position of the stream except position 1st, 2nd, 9th and 10th. The duration of this session was approximately 25 minutes.

Repetition Task

Each of the 18 blocks of the Repetition Task consisted of 11 trials. The trials started with the instruction "Search for the repeated word". Then participants were presented with 11 trials with a similar structure as those of the Detection task. However, in this task, participants were instructed that one of the words presented during the first trial would be presented several times in the following trials. This first trial of the block was not followed by a question[3], while the remaining 10 trials were followed by the question "Have you seen the repeated word?". Participants typed the word that they considered was repeated or pressed "Enter" to continue with the next trial. As in the Detection task, the repeated item (RI) was presented in 5 of the 10 experimental trials, so that half of the experimental trials in a block did not contain the RI. Importantly, distractors did not repeat within a block. The administration of this task lasted approximately 28 minutes.

2.2 Experiment 2

The materials and procedure were the same as for Experiment 1. The only difference was that in this experiment a different set of SOAs was selected: 17, 50, 84, 117, 250 and 400ms.

[3] Note that in Bowman et al., (2014) participants were not instructed of the presence of the repeated name at a specific trial, making the task harder. This may explain why the repeating stimulus was detected less easily in Bowman et al., (2014).

Participants

24 undergraduate students from the University of Birmingham participated in Experiment 2 in exchange for course credits. None of these students had participated in Experiment 1. The same exclusion criteria were used in both experiments. The experiments performed conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Results

Repetition and Detection Task

Although the main piece of evidence we are seeking to identify is the estimation of the (accumulation) repetition effect in RSVP (which is described in the next section), the results described here –comparing the d' score of the Detection and Repetition task– give a global picture of the time-course of participants' performance in both task.

Experiment 1

We calculated participants' d' measures for both tasks. In the Detection task, the Hits, Misses, False Alarms and Correct Rejections were obtained from the "yes" and "no" responses. In the case of the Repetition Task, we considered as Hits all target-present trials where participants typed any word, even if it was not the Target. Similarly, we considered as False Alarms all trials where participants typed a word at the end of the trial but the repeated word was not presented. It is important to notice that the same criterion (any word, correct or incorrect, was typed at the end of the trial) was used for the classification of responses as Hits or False Alarms[4]. Therefore, participants' d' measures above zero will correspond to their ability to distinguish trials where the repeated word was present from trials without the repeated word.

A 2 (Task: Detection and Repetition) x 6 (SOA: 33, 133, 233, 333, 433, 533) within subjects ANOVA revealed that the d' values (see Figure 2 and Table 1) were significantly larger for the Detection Task (mean=2.87 and sd=1.52) than for the Repetition task (mean= 1.51 and sd=1.46), $F(1,20)=171.7$, $MSE=115.78$, $p<0.001$. There was also a main effect of SOA, $F(5,100)=91.93$, $MSE=32.85$, $p<0.001$. Unsurprisingly, d' values were larger for the long SOA conditions, see Table 1. The interaction was also significant $F(5,100)=7.75$, $MSE=1.97$, $p<0.001$, showing the task effect was not constant across the SOA conditions. To analyse the interaction, we subtracted the d' values of both task and performed post-hoc comparisons across the SOA levels. Pairwise comparisons revealed that the task effect was larger at 133ms relative to 33, 333, 433, and 533ms (all $ps<0.05$, Holm corrected). Similarly, the task effect at 233ms was significantly larger than the task effect at 33ms ($p<0.05$, Holm corrected).

Experiment 2

A 2 (Task: Detection and Repetition) x 6 (SOA: 17, 50, 84, 117, 250, 400) within subjects ANOVA showed a main effect of Task, $F(1,23)=247.2$, $MSE=97.09$, $p<0.001$

[4] The consequence of this is that any "incorrect-identity" repetition-seen response that could be made is truly an error response and is thus equally likely to arise in a target-present as a target-absent trial and is as likely to increase False Alarms as Hits. This will then just show up as a change in response-bias, but will not impact d-prime, which is how it should work in signal-detection theory. Figure A.1.2 in Appendix 1.5 shows performance in the repetition task using an alternative criterion: any string with a Levenshtein distance (the minimum number of single-character edits; insertions, deletions or substitutions) of 2 or less was accepted as a correct response.

(Detection task: mean= 2.40, sd=1.9; Repetition task: mean=1.24, sd=1.58). The main effect of SOA was also significant, F(5, 115)= 164.4, MSE=61.56, p<0.001. Finally, the interaction Task x SOA was significant F(5, 115)= 15.17, MSE=4.92, p<0.001 (see Figure 2 and Table 2).

Pairwise comparisons of the task effect (d' Detection task – d' Repetition task) across SOA conditions revealed that the task effect was significantly larger (all ps<0.05, Holm corrected) at 117ms relative to all other SOA conditions (17, 50, 250 and 400ms) except the 84ms condition (p=0.67, Holm corrected). In contrast, the task effect at 17ms was significantly smaller (all ps<0.05, Holm corrected) compared with any other SOA condition (50, 84, 117, 250 and 400ms).

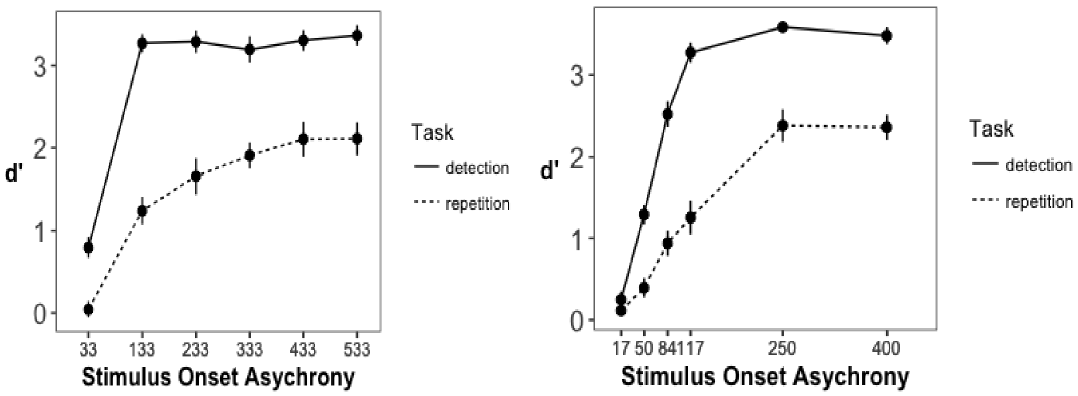

Figure 2. Results for experiments 1 (left) and 2 (right). The graph shows the d' scores for the Detection and Repetition tasks across the SOA conditions. Error bars indicate standard errors of the mean.

TABLE 1. MEAN D' VALUES AND (STANDARD DEVIATIONS). EXPERIMENT 1

TASK	SOA					
	33	133	233	333	433	533
DETECTION	0.79 (0.57)	3.27 (0.51)	3.28 (0.62)	3.19 (0.60)	3.30 (0.58)	3.36 (0.56)
REPETITION	0.04 (0.44)	1.24 (0.76)	1.65 (1.02)	1.91 (0.71)	2.11 (0.97)	2.11 (0.91)

TABLE 2. MEAN D' VALUES AND (STANDARD DEVIATIONS). EXPERIMENT 2

TASK	SOA					
	17	50	84	117	250	400
DETECTION	0.25 (0.51)	1.29 (0.59)	2.52 (0.77)	3.27 (0.60)	3.59 (0.31)	3.48 (0.5)
REPETITION	0.12 (0.30)	0.39 (0.57)	0.94 (0.75)	1.25 (1)	2.38 (0.98)	2.36 (0.74)

The effect of repetition in RSVP

In this section, we assess the repetition effect on the probability of identifying a repeated item (RI) for the first time in the Repetition task. Specifically, we tested whether the probability of identifying the RI increased as a function of how many times the RI had been presented. In the Repetition Task, participants were presented the RI five times (1st, 2nd, 3rd, 4th and 5th repetition) per block. This allowed us to calculate the probability of identifying the target at each repetition for the first time. Note that, contrary to the procedure adopted by the analysis described above (where any string of letters typed by participants was included in the calculation of Hits and False Alarms) for the present analysis, we only considered responses as correct if the RI was typed. To do that, the number of words correctly identified for the first time at repetition j (1 to 5) was divided by the number of blocks where the RI was missed at every repetition before j. That is, at repetition 1, the number of RIs identified (as repeating) at that repetition was divided by the total number of blocks (3) at each SOA condition. For repetition 2, the number of RIs identified (as repeating) for the first time at repetition 2 was divided by the total number of blocks where the RI was missed at repetition 1, and so on. See Appendix 1 for a formal definition of this measure.

To calculate this probability at every repetition it is essential that participants missed some repetitions, since the first-seen measure depends on the number of repetitions where the target was not detected. From the results of the repetition task (see Figure 2, it is evident that participants were very good at detecting repetitions at long SOAs (>200ms). They had identified correctly on average more than 50% of the words by the third repetition. Consequently, we considered that the data available at these long SOA conditions did not give enough power to provide a good estimation of the effect of repetition. The opposite pattern was found at the shortest SOAs (33ms, Experiment 1; 17ms, Experiment 2) where participants were effectively at floor. Therefore, the data submitted to statistical analysis were those corresponding to the SOA conditions: 133ms in Exp. 1, and 50, 84 and 117 in Exp.2). For completeness of reporting, Appendix 1.3, Tables A.1.2 and A.1.3 summarise the full set of results.

Experiment 1

The data (see Table 3) were analysed using a probit mixed-effect model, with the lme4 package in R (Bates, Mächler, Bolker, & Walker, 2015). Repetition (1^{st} to 5^{th}) was included as a fixed factor. The random structure of the model only included the intercept of the Participant factor. More complex models (including those with slopes in the random factor) did not reach convergence. A Chi-square test from the probit linear mixed model results showed a non-significant effect of repetition ($p > 0.9$).

The non-significant effect of the repetition indicates that we cannot reject the Null Hypothesis that there is no effect of Repetition. To assess the extent in which our data are more consistent with the Null Hypothesis, we calculated the Bayes Factor from the Bayesian Information Criterion (BIC; Raftery, 1995) of two models following the approximation suggested by Wagenmakers (2007)[5]: the full model (Repetition + random intercept per participant) and the restricted model (Intercept + random intercept per participant). The Repetition factor was included as a linear regressor (1 to 5) to model the "build-up" mechanism proposed by Endress and Potter (2014). The BF_{01} (null/alternative) was 6.90. Following the classification scheme of Raftery (1995) we concluded that this is positive evidence (BF 3-20) for the Null Hypothesis of no Repetition effect.

Experiment 2

---

[5] exp((BIC_1 - BIC_0)/2)

We analyzed the data from Experiment 2 using the same methodology as in Experiment1. The data of the probability of first identification as a repetition (see Figure 3 and Table 4) were analyzed with a probit linear mixed model with SOA (50, 84, 117) and Repetition (1st to 5th) as fixed factors. The random structure of the model included the intercept per participant. The Chi-square test from the model revealed a significant effect of SOA, $\chi2(3) = 7.46$, $p = 0.024$, and no significant effect of Repetition nor the interaction SOA x Repetition (both ps>0.8).

The Bayes Factor between the BIC of the full model (SOA, Repetition and SOA x Repetition) and the restricted model (SOA) showed that the restricted model was favored by a factor of 9.077142e+13. Finally, the BF for each SOA condition between the full model (Repetition + random intercept of the participants) and the Null (intercept only + random intercept of the participants) revealed the following results: BF_{01} SOA 50ms = 8.09, BF_{01} SOA 84ms = 10.18, BF_{01} SOA 117ms = 9.99. As from experiment 1, the Bayes Factor analysis indicates that the results from experiment 2 favored the Null Hypothesis of an absence of a Repetition effect.

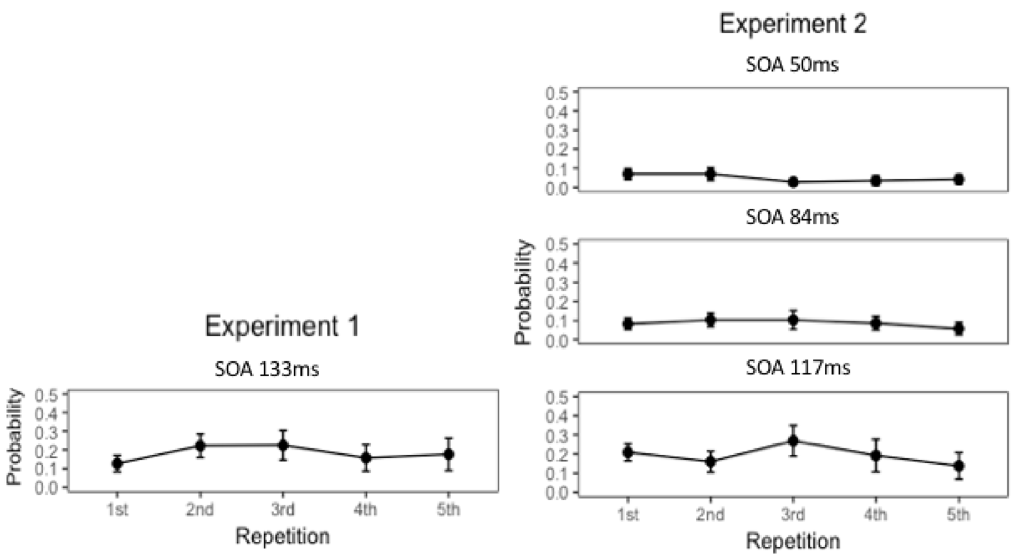

Figure 3. Results for experiments 1 (left) and 2 (right). The graph shows the probability of first detecting the target (as repeating) after every repetition. Error bars indicate standard errors of the mean.

Table 3. Mean probability of first detection as a rep. and (standard deviations). Experiment 1

	Repetition				
SOA	1	2	3	4	5
133	0.13(0.2)	0.22(0.29)	0.23(0.35)	0.16(0.3)	0.18(0.36)

Table 4. Mean probability of first detection as a rep. and (standard deviations). Experiment 2

	Repetition				
SOA	1	2	3	4	5
50	0.07(0.14)	0.07(0.17)	0.03(0.1)	0.03(0.12)	0.04(0.11)
84	0.08(0.15)	0.1(0.17)	0.1(0.24)	0.09(0.18)	0.06(0.16)

117	0.21(0.22)	0.16(0.27)	0.27(0.38)	0.19(0.38)	0.14(0.29)

3. Discussion

In two studies, participants were presented streams of words at different presentation rates, while performing a detection or a repetition task. The repetition task required participants to search for repeated words across several streams (different trials). The detection task instructed participants to search for pre-specified target words within each trial. Participants were able to complete both tasks, with performance being considerably better in the easier detection task.. More importantly, in the repetition task the probability of detecting a repetition for the first time did not increase with the number of repetitions. This suggests that the performance in the repetition task was not aided by evidence accumulation across repeated instances of the same word. This result has implications for our understanding of how fleeting representations are processed and remembered. At SOAs in the range of 84-133 ms, participants exhibit excellent performance on the detection task, with d' scores well above 2.0. Thus, it must be the case that most of the words are individually being processed to some degree, and yet at these same rates, repetition detection does not benefit from evidence accumulation.

In previous studies, the absence of accumulation effects had been obtained in conditions where performance in memory tasks was extremely poor. Subramanian et al., (2000) showed that non-targets repeated up to 15 times were not detected better once they became targets in a detection task. Importantly, they reported that, in similar conditions, recognition for non-targets in a forced-choice task was at chance. In Bowman et al., (2014), participants found it very hard to detect repeated names —presented up to 50 times— even when they were instructed to search for them. These findings pointed to a failure of memory encoding for items in RSVP. It seems that items neither accumulated evidence nor were easily retrieved, since there was little evidence for memory representations that could carry out those functions. The present results agree with those of (Bowman et al., 2014) and Subramanian et al., (2000) in that no repetition accumulation effect was observed. However, contrary to these studies, participants' performance in the repetition task was not at chance. That is, even though participants' capacity to notice repetitions was somewhat spared, the repetition accumulation effect was absent.

In contrast with the present results, in other studies, repetition demonstrably improved recognition performance, even in RSVP experiments. For example, Endress and Potter (2014) found that items presented more often in RSVP trials were recognized more accurately. They suggested that the repeated instances of items in RSVP build-up and result in stable long term memory representation. Our results could be signalling one fundamental limit of this mechanism. While it might be true that repetition facilitates retrieval, it could be that the retrievability (\ conscious perception) of items is a necessary condition for any repetition accumulation effect, as per case 1 of figure 1. Alternatively, it would have been possible that memory representations could accumulate evidence through repeated exposure prior to being strong enough to be consciously perceived and explicitly retrieved, i.e. case 2 of figure 1. Our findings contradict this notion, at least in experimental paradigms such as RSVP. Our results showed no repetition accumulation effect prior to the explicit recognition of words (as repeating) for the first time, i.e. case 3 of figure 1. This opens the possibility that in RSVP, the evidence accumulation process is restricted to those items that have been consciously perceived and thus encoded in strong memory representations, at least sufficiently strong to support explicit recognition. In fact, the findings in Endress and Potter (2014) are consistent with this possibility. Since the benefit of repetitions was only apparent in a final recognition test (by comparing accuracy for items presented less and more often), it could be that only those items consciously detected were those that improved with repetitions after that point.

Alternatively, the absence of a repetition effect could be attributed to memory capacity limits. Some stimuli would not be encoded in memory since memory limits are reached –given that streams consisted of 10 words (above short-term memory span). However, the d´ results of the Repetition Task contradict this interpretation. As can be observed in Figure 2, performance steadily improves up to asymptotic levels at slow presentation rates (SOA >200ms). If performance in the Repetition Task were to be explained only by memory capacity limits, no effect of presentation rate would be expected. Instead, SOA strongly modulated the results, suggesting that the presentation conditions (duration on screen and masking by subsequent items) played a crucial role. Taken together, the results of the Repetition Task seem to indicate that, in respect to memory encoding, items in RSVP could result in two qualitatively different types of representations. Some items could be encoded in sufficiently strong memory representations that successful explicit memory recognition is possible. On the other hand, most items may only elicit extremely fleeting "percepts" that would not survive the presentation of new items.

The present results do not contradict the notion that the presentation of items in RSVP can facilitate/prime the processing of subsequent stimuli, which can be considered an alternative form of evidence accumulation. Priming effects are very unlikely to be observed in a paradigm such as the Repetition Task. The distance between RIs –both in terms of number of intervening stimuli and duration– is relatively long and priming effects tend to be inversely related to the gap between items of interest (Ferrand, 1996).

Alternatively, priming might have an effect in a recognition task (and the Repetition Task can be conceived as a continuous recognition task) by increasing the familiarity of repeated items. However, it is unlikely that our Repetition Task could be performed on the basis of familiarity decisions given that the format of the task, where an explicit report of the identity of the RI is required, makes the task similar to paradigms (such as recall), which rely on recollection processes – and not familiarity.

Are unidentified instances of RIs below threshold? The definitive answer to this question might ultimately require a resolution of the ongoing debate about the existence of phenomenological consciousness (Block, 2007; Cohen & Dennett, 2011). Indeed, we cannot rule out the possibility of phenomenological awareness of items presented in RSVP, i.e. that they are perceived, but forgotten before the end of the stream, see Figure 4. However, the most commonly accepted view in the RSVP literature is that the report of items at the end of the stream is taken as the behavioural correlate of conscious perception (for example: failure to report items in the attentional blink window) (Bergström & Eriksson, 2014). In the present experiment, the missed repetitions of RIs could similarly be interpreted as instances of below threshold stimuli.

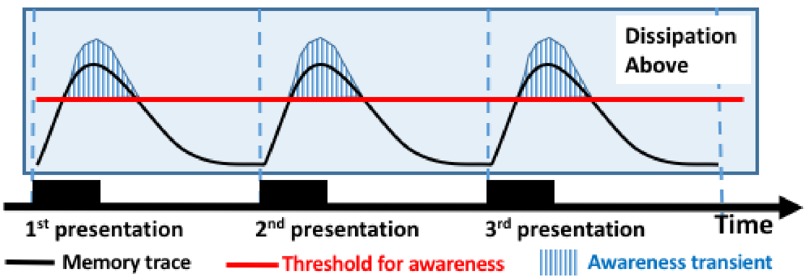

Figure 4. Further possibility for how the brain responds to repeated presentations. The stimulus sequence is shown in black as three presentations of the same stimulus. The awareness transient reflects the conscious experience of the presented stimulus. Evidence dissipates between presentations, all of which generate a conscious percept.

This said, even if one accepts the contribution of phenomenological awareness to perception in RSVP, i.e. that items in this presentation format could be consciously perceived, but then rapidly forgotten, our line of reasoning is not in fact contradicted. That is, we have demonstrated a failure of evidence accumulation of repeated presentations, and if some of these “prior” presentations elicited above threshold experiences that were then forgotten, that would just represent evidence for the even stronger claim, that it is also possible that evidence does not accumulate for items that are consciously perceived. Importantly, though, it would seem highly unlikely that every prior presentation of the RI is consciously perceived and then forgotten. Our key findings are based upon streams with SOAs between 50 and 133ms; it seems inconceivable that every item in a 10 item stream at such SOAs would be consciously perceived. Accordingly, even if some RIs are consciously perceived and then forgotten, there remain some that would not, whose existence is sufficient to carry our claim.

The absence of apparent evidence accumulation (for stimuli in RSVP) could be the consequence of encoding processing limits of a similar nature to those responsible for other empirical findings such as the Attentional Blink (Raymond, Shapiro, & Arnell, 1992) and Repetition Blindness (Kanwisher, 1991). 2-stage theories of temporal attention (Bowman & Wyble, 2007; Chun & Potter, 1995) propose a first stage of large capacity where the attributes of items are extracted and a second stage where items are consolidated in WM. The memory encoding of items in our experiments seems to follow a 2-stage behaviour, where only a few items would break through to a second stage of consolidation into a durable representation (Wyble, Bowman, & Nieuwenstein, 2009). Indeed, the encoding of stimuli may depend on (what could be considered) random factors associated with presentation and processing (e.g. how effective the masking induced by the following item). This 2-stage interpretation gains further support if we compare the results of our repetition task with those of the detection task. Participants are much better at detecting pre-specified task-relevant items at any SOA condition than searching for repetitions. As can be observed in figure 2, at presentation rates close to 10/s, participants are almost at ceiling at discriminating target-present from target-absent streams. This is consistent with (Potter, 1976), and what was called (Sub)liminal Salience Search in (Bowman et al., 2013). It also fits with the idea of featural and semantic information being extracted in a first large-capacity stage; and a natural prediction then is that most items could be detected on the basis of their task relevant properties, which is in agreement with the results of the Detection Task.

The findings also fit with a stronger claim, viz, that a key aspect of conscious perception is supporting *episodic* encoding of items, a position that Kanwisher has also argued for (Kanwisher, 2001). One interpretation of the repetition task is that it is episodic in nature: to know that a previous instance has occurred, an individuated representation of that earlier occurrence needs to be represented in memory. This is exactly the role of tokens in the Simultaneous Type/ Serial Token model (Bowman & Wyble, 2007): when bound, they indicate the occurrence of a type. This system detects a repetition when the same type is bound to two tokens. Thus, tokens provide durable representations of event occurrences. If the preconscious percept does not induce a durable representation, it certainly cannot provide a tokenized (episodic) representation. This leaves the possibility that our subconscious is, strictly speaking, *episodically blind* and even potentially that the very process of perceiving involves episodically tagging experiences. This is what we call the Tokenized Percept Hypothesis, from which we will seek in future experiments to find further evidence.

In summary, we have shown an effect of explicit recognition of items in the absence of any benefit obtained from successive repetition. These results suggest that the process of episodic encoding in RSVP follows a bottleneck behaviour, where most items fail to be durably consolidated in memory or even to leave any memory trace that is able to support evidence accumulation. Combined with the results of the detection task, the findings suggest that in fast RSVP streams, items can be discriminated easily on the basis of their perceptual attributes in conditions where only a small subset of them can be durably encoded. Taken together, these results agree with 2-stage theories of temporal attention, where a first stage of large capacity consists of the activation of the perceptual properties of items (enabling, for example, effective detection performance) and a second stage encodes a subset of items into durable WM. Insofar as it is assumed that conscious percepts are those that can be reported/retrieved, these results suggest that conscious perception is required for an additional function: the accumulation of evidence through repetition.

This chapter has provided evidence indicating the fragility of episodic encoding in Rapid Serial Visual Presentation (RSVP) tasks. However, this evidence was primarily derived from investigating the influence of previous item presentations on subsequent repetitions. In Chapter 3, we will directly address the matter of item reportability in RSVP by examining performance in free-recall and recognition tasks.

Chapter 3: Fragile memories for fleeting percepts

Bowman, H., & Avilés, A. (2021, March 31). Fragile Memories for Fleeting Percepts. https://doi.org/10.31234/osf.io/fr3bg

1. Introduction

Our brains are strikingly good at finding stimuli of interest to us in our environment. This is compellingly demonstrated by detection and identification performance for targets presented in RSVP (Rapid Serial Visual Presentation) streams. For example, accuracies of 80% and above are often reported when identifying a single letter presented within an RSVP stream of digits, e.g. T1 performance @ lag-8 in Craston et al. (2009)); also see, (Potter, 1976; Potter et al, 2014; Barnard et al, 2004; Bowman et al, 2013).

This, then, suggests that the brain is performing a *search* when viewing RSVP, i.e. it is attempting to “find” salient stimuli, which are made available to conscious experience once found. A key question is the extent to which this search is selective. Is the brain really picking out just the salient stimuli at the expense of the vast majority of the other stimuli in the stream? Does it have the capacity to be so selective? This is important to understand, since if there is very little trace left in the brain for the stimuli not selected, it would suggest that the searching for salient stimuli is largely a subliminal process, i.e. non-salient stimuli are rejected without them having to be consciously processed.

So, a key question to answer is how many RSVP items need to be deeply processed to obtain the level of detection and identification performance observed. In particular, since the position in a stream in which targets are presented is not predictable, an 82% identification performance would seem to suggest that ~80% of items in the stream need to be sufficiently deeply processed that they can be either accepted as a target or rejected as a distractor. This is a very large proportion of the stimuli presented in the stream. However, what does "sufficiently deeply" mean?

One way to assess this is to determine the memory traces left by arbitrary stimuli in streams. High identification performance in RSVP makes clear that *target* stimuli typically leave strong memory traces. What, though, of the memory traces left by *non-targets*, i.e. task-irrelevant distractors? If the process of rejecting non-targets as not of interest was performed consciously, a relatively robust memory trace could be expected to form. This is the question we investigate in this paper. To do this, we build from the seminal work of Mary Potter (Potter, 1976), who showed a substantial difference between detection and recognition performance; see figure 1.

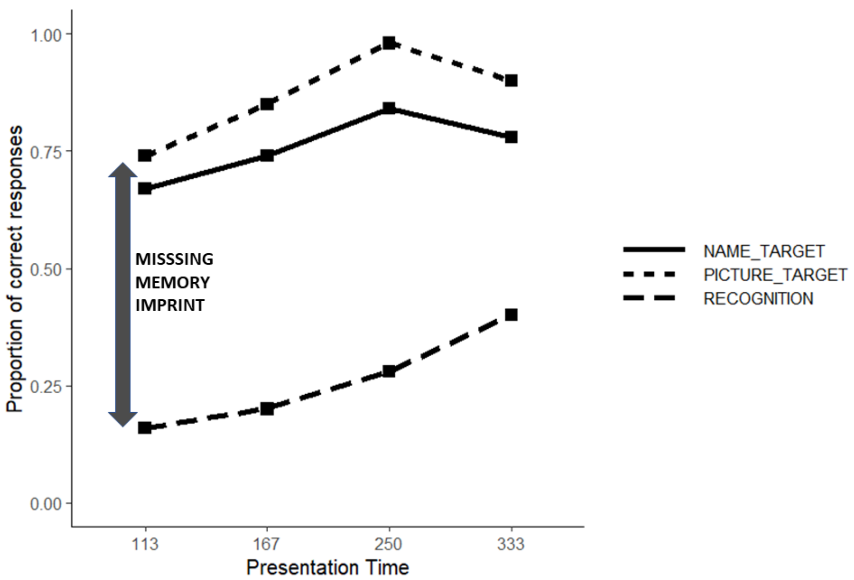

Figure 1: first finding of a substantial disparity between detection performance (picture target and name target lines) and memory (recognition line) at rapid serial visual presentation rates (113ms and 167ms) and slower (250ms and 333ms). Strikingly, at the fastest rate, detection performance is over four times recognition memory, indicated with

added annotation as missing memory imprint. Figure adapted from Potter, M. C. (1976). Short-term conceptual memory for pictures. Journal of experimental psychology: human learning and memory, 2(5), 509.

A first objective of our experiments was to refine the approach of Potter (1976) to the particular question we are exploring. Specifically, Potter (1976) tested recognition memory for items arising throughout the course of the RSVP stream. This, on average, gives a considerable period of time for items to be forgotten between presentation and recognition test. As a result, Potter (1976) will underestimate the memory trace for a stimulus that may exist soon after its presentation. In response, we specifically probe memory for the *last* three items presented before an end of stream mask. This gives as much chance as possible for a strong memory to be observed if it exists.

Associated with this refinement, we also randomly varied the length of streams, in order to prevent participants from being able to build-up an expectation for when the end of stream will happen and differentially focus on remembering the last three items. This contrasts with Potter (1976), where all streams were of the same length.

Recency effects are a key indicator of the fragility of memory representations (Broadbent & Broadbent, 1981). That is, extreme recency effects, where, say, just the last item presented is reliably remembered, would suggest a rapid dissipation of memories, and marked fragility. Accordingly, we look for increasing memory across the last three items in RSVP streams.

Additionally, the nature and depth of memory encoding during RSVP streams can be assessed by comparing *free recall* with *recognition.* Free recall tests assess memory representations that are freely accessible, which one would expect to be currently present in working memory. In contrast, recognition memory tests can be performed on the basis of mere familiarity, i.e. a sense that the probed stimulus has been seen previously, without a full source memory of where or when the stimulus was experienced.

These issues provide important ways to test the Simultaneous Type/Serial Token (STST) model of temporal attention (Bowman et al., 2008; Bowman & Wyble, 2007) as a theory of conscious perception, and the associated tokenized-percept hypothesis (Chapters 1 and 2). Importantly, task-irrelevant distractors are only represented in the first stage of the STST model, and therefore do not reach working memory, which resides in stage two. This suggests the following hypotheses about the memory traces of task-irrelevant distractors (i.e. non-targets) presented in RSVP,

1) they should decay rapidly;

2) recognition performance should be substantially better than recall performance, since familiarity can be determined from first stage representations;

3) recency of recognition memory should be stronger than recency for recall. These hypotheses also resonate with the idea that RSVP streams provide *glimpse percepts*, as suggested by the Glance-Look model of the attentional blink (Barnard & Bowman, 2004; Su et al., 2011).

To assess the impact that conscious perception has on the fragility of memories, we explore three Stimulus Onset Asynchronies (SOAs): 117ms, 230ms and 350ms. The first of these is a typical RSVP rate, with stimuli presented "on the fringe of awareness", or, indeed, below it, while the last should be a largely conscious regime, in which the vast majority of stimuli would generate a conscious experience.

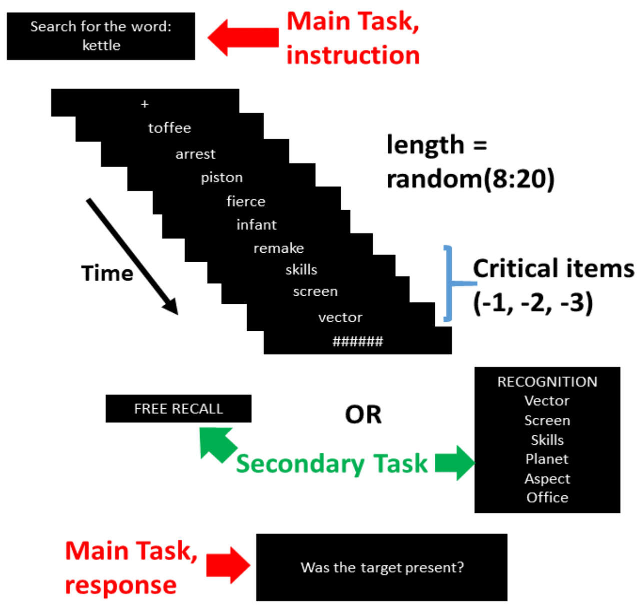

Figure 2: experimental procedure: RSVP streams of words were presented, terminated by a mask. Streams were stopped randomly from the 8th to the 20th item. The main task was to detect the specified word and perform a presence/ absence judgement. Additionally, in a secondary task, we either asked participants to freely recall words from the stream, or to perform a recognition test, which included the last three words in the stream (positions -1, -2 and -3), randomly placed amongst three unpresented words.

2. Method

Participants

26 (mean age= 19.7, sd= 1.2) undergraduate students (5 males) of the University of Birmingham took part in the experiment in exchange for course credits, and gave informed consent. All were native English speakers and had normal or corrected-to-normal vision. Sample size was estimated using GPower (Erdfelder et al., 1996). The projected sample size was obtained relative to the effect of Position in Recognition, η_p^2=0.1 (based on data from pilot studies). With an alpha = .05 and power =

0.90, the projected sample size is approximately N = 20.52 (one-way Anova with 3 repeated measurements). Our sample size is slightly bigger than the projected sample size and should be adequate for the main objective of the study. We based the power analysis on the effect of Position in Recognition since that was smaller than the effect of SOA in Recognition (based on the pilot). The effects of Position or SOA in Free-recall were not considered for the power analysis since we did not have pilot data on these effects and we failed to find similar studies in the RSVP literature. The replication study (see Appendix 2) confirmed that a sample size close to that used in the main experiment was adequate to replicate the findings. The experiment conformed to British Psychological Society criteria for the ethical conduct of research. All experimental protocols were approved by the Ethics Committee at the University of Birmingham.

Materials

1700 English words were selected from (Warriner et al., 2013). From this main list, 48 words were selected randomly to serve as Targets in the Detection task for each participant. Additionally, a set of 72 words were selected to serve as New Words in the Recognition task. All words had 6 letters and none of them were proper nouns.

Procedure

RSVP streams were presented on a 24’’ LCD screen (refresh rate: 60Hz, resolution: 1600 x 1200) using custom PsychToolbox (Brainard, 1997; Kleiner, 2010; Pelli, 1997) scripts running under Matlab 2016a. Stimuli were 16-point white (75% white) characters on a dark background (25% white). Participants were seated at 60cm from the screen. The experiment was conducted individually in a quiet room. Each experimental session consisted of 96 trials (32 trials per SOA condition: 117, 230, and 350ms).

The format of the experiment is shown in figure 2. At the beginning of each trial, participants saw the instruction “Search for the word:” in the upper part of the screen along with the target word presented at the centre of the screen. They were instructed to press the “space bar” after reading the target word to continue. Each stream started with the centred presentation of the fixation cross (+) for 500ms. Then a sequence of words was presented and a final item “######” ended the stream. Crucially, the length of the streams varied randomly between 8 and 20 words in each trial.

The trial ended with the presentation of one or two successive questions — (the first one) presented immediately after the final “######” item. There were four types of trial, each comprising 25% of the total trials, which were randomly intermixed. These four trial types were as follows.

a) Detection trials: Participants were presented with the question: “Have you seen the target?” They pressed the “left arrow key” to indicate YES or the “right arrow” to indicate NO.
b) Recognition + detection trials: the question “Which of the following words were presented?” was displayed in the upper part of the screen. A column with 6 words were presented, each one preceded by a number from 1 to 6. Half of the words corresponded to the items presented in the streams at positions -1, -2, and -3. The other 3 probes were new words. The order of presentation of the 6 probes was random. Participants were instructed to press the keys with the number of the probes (1 to 6) that they considered as old words. Then they pressed “Enter” to move to the Detection question, which was identical to that of the Detection only trials.
c) Free-recall + detection: the instruction “Type as many words as you can remember” was presented. The words typed by the participants appeared at the

centre of the screen. They were instructed to press “Enter” when they had typed all words they remembered or in case they could not remember any words. After this, they were presented the same Detection question as in the Detection only trials.

d) Filler task: the question “How many words have you seen?” was presented. They responded by pressing the keys that corresponded to the number that they estimated. After pressing “Enter”, participants were presented with the Detection question. The results of this task are not reported here.

Results

Detection task

We calculated the d’ score from the Hit and False Alarm rates obtained in the responses to the Detection question. A one-way (SOA: 117, 230, and 350ms) within-subjects ANOVA showed significant differences on the d’ scores, $F(2,50)=3.50$, $MSE =0.98$, $p<0.05$, $\eta_p^2=0.12$. d’ values were slightly larger for the longer SOA conditions (SOA 230: mean=2.57 and sd=0.63; SOA 350: mean=2.45 and sd=0.62) than for the short SOA (SOA 117: mean=2.2 and sd=0.61).

Recognition

Hits and FAs rates were calculated from the responses to the recognition question. Hits were obtained for stimuli presented in positions -1, -2, and -3 in the stream. We considered as Hits yes responses to old-word probes. False alarms were yes responses to new-word probes. A 3 (SOA: 117, 230, 350) x 3 (Position: -1, -2, -3) within subjects ANOVA revealed that the d' values (see Figure 3 and Table 1) were significantly larger for the longer SOAs, $F(2,50)=17.97$, $MSE=10.22$, $p<0.001$, $\eta_p^2=0.43$, and for words presented later in the stream, $F(2,50)=11.77$, $MSE=2.48$, $p<0.001$, $\eta_p^2=0.16$. The interaction was not significant $F(4,100)=0.5$, $MSE=0.13$, $p>0.1$. This pattern of results did not vary depending on the responses to the Detection question. That is, the results were similar when just Hit or Correct Rejection trials were included in the analysis.

Free Recall

We calculated the proportion of correct responses (see Figure 3 and Table 1) from the number of times participants included the stimuli presented in positions -1, -2, or/and -3 in their responses to the free-recall question. The 3 (SOA: 117, 230, 350) x 3 (Position: -1, -2, -3) within subjects ANOVA showed significantly higher accuracies for longer SOAs, $F(2,50)=14.16$, $MSE=0.39$, $p<0.001$, $\eta_p^2=0.31$, and a significant effect of position, $F(2,50)=27.72$, $MSE=0.81$, $p<0.001$, $\eta_p^2=0.48$. The interaction did not reach significance, $F(4,100)=0.86$, $MSE=0.02$, $p>0.1$. Again, the same pattern was found when just Hit or Correct Rejection (for Detection) trials were considered.

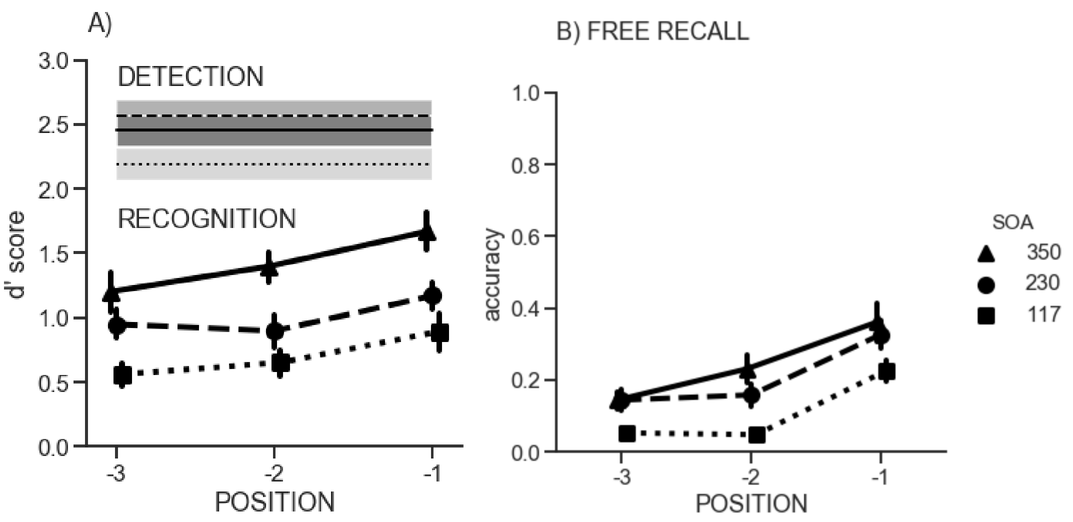

Figure 3. A) Recognition (lower nine points) and detection (upper three lines) performance (left), for three SOAs: solid line 350ms, dashed 230ms, dotted 117ms. Only recognition performance is able to vary with position relative to end of stream: -1, -2, -3. B) Recall performance (right) for the three SOAs and three positions. Error bars and shadowed area indicate standard errors of the mean.

Table 1. Mean d' values and (standard deviations) for the Recognition task and mean correct responses and (standard deviations) for the Free-recall task

	Recognition			Free-recall		
SOA	117	230	350	117	230	350
POSITION						
-1	0.89	1.17	1.67	0.23	0.33	0.36
-2	0.65	0.9	1.4	0.05	0.16	0.23
-3	0.56	0.95	1.2	0.05	0.14	0.14

Detection vs. Recognition

Since detection cannot change with position (-1, -2, -3), we compare detection and recognition with position collapsed, giving a 2 (Task: detection, recognition) x 3 (SOA: 117, 230, 350) within-subjects ANOVA. This showed that d' values were higher for Detection (mean= 2.4, sd=0.63) than for Recognition (mean=1.04, sd=0.57), $F(1,25)=208.9$, $MSE =71.93$, $p<0.001$, $\eta_p^2=0.88$. As expected, the main effect of SOA was also significant, $F(1,25)=12.16$, $MSE =3.3$, $p<0.001$, $\eta_p^2=0.4$. Furthermore, the interaction also reached significance, $F(2,50)=5.5$, $MSE =1.1$, $p<0.05$, $\eta_p^2=0.18$, which suggests that the effect of SOA was higher in the Recognition task relative to Detection. (The potential confound that detection is at ceiling is discussed and ruled out in the appendix 2.1.)

Recency effects

In order to compare recency for recognition and recall we operationalise recency as the ratio shown in figure 4, which we call *normalised recency*. With this transformation, we aim to show the contribution of the responses for each position (-1, -2, -3) to the total performance in the task, where total performance is summed performance across all positions. For example, the largest recency effect would involve a normalised recency of 1 at the -1 position, reflecting that performance (d' or accuracy) is zero at -3 and -2 and so the performance at -1 *is* the total performance. Additionally, by turning our dependent measure into a proportion (i.e. of total performance), we can directly compare the recency patterns for Recognition and Recall, which we could not with d' and accuracy (see Table 2).

A 2 (Task: free-recall, recognition) x 3 (SOA: 117, 230, 350) x 3 (Position: -1,-2,-3) within-subjects ANOVA revealed that the position main effect was significant,

$F(2,44)=49.77$, $MSE=2.97$, $p<0.001$, $\eta_p^2=0.55$. Position interacted significantly with task $F(3,69)=12.45$, $MSE=0.68$, $p<0.001$, $\eta_p^2=0.3$ and with SOA $F(6,144)=3.11$, $MSE=0.13$, $p<0.05$, $\eta_p^2=0.14$. Finally, the three way interaction was significant $F(6,135)=2.65$, $MSE=0.1$, $p<0.05$, $\eta_p^2=0.11$. We excluded the main effects of Task and SOA from the model, since it was singular if either or both were present. This is essentially because the ratios we take for our normalised recency measure, imply that, by construction, there are no differences between levels of these main effects (for example, the sum of all Recognition data points is 3 and the sum of all Recall data points is 3; see figure 4).

We explored the three way interaction by running independent 3 SOA x 3 Position within-subject ANOVAs for each task, which are the simple-effects of the three-way interaction. The results showed a significant interaction in the Free-recall task, $F(6,135)=4.15$, $MSE=0.21$, $p<0.001$, $\eta_p^2=0.16$. The same interaction did not reach significance in the Recognition task ($p>0.1$).

A) Recency operationalised as:

$$\forall j \in \mathbb{Z}(-3 \le j \le -1) \cdot norm\_rec_{(j)}(X) = \frac{X(j)}{\sum_{i=1}^{3} X(-i)}$$, with X a performance measure

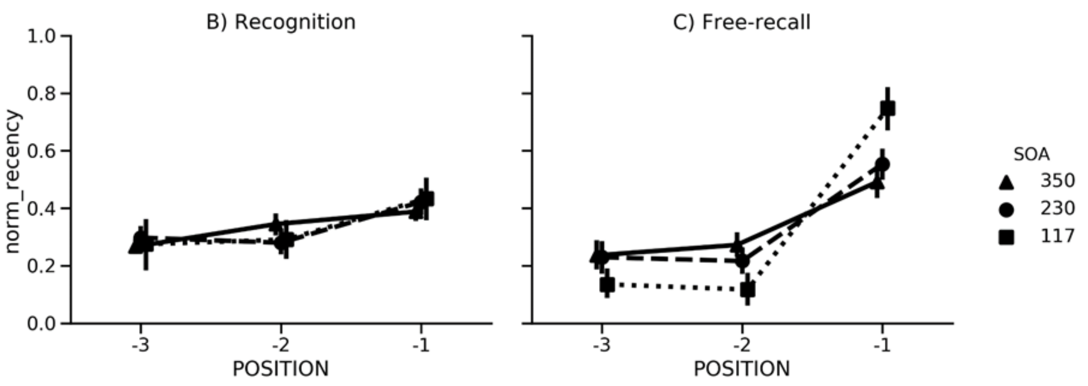

Figure 4: Recency results for three SOAs and the last three stream positions: A) formula for normalised recency, expressing performance at each position relative to total performance across all positions. *norm_rec(-1)* is the proportion of total performance that

is seen at the last position of the stream; it most directly characterises recency. B) Normalised recency for recognition, with performance measured as d-prime, showing a recency effect that is essentially the same for all SOAs. C) Normalised recency for recall, with performance measured as accuracy, showing a substantially larger recency effect for a 117ms SOA.

Table 2. Mean "normalized" values (proportion in which each position contributes to total performance) and (standard deviations) for the Recognition and the Free-recall tasks.

	Recognition			Free-recall		
SOA	117	230	350	117	230	350
POSITION						
-1	0.43 (0.36)	0.42 (0.19)	0.39 (0.13)	0.75 (0.32)	0.55 (0.24)	0.49 (0.24)
-2	0.29 (0.32)	0.28 (0.17)	0.34 (0.16)	0.12 (0.24)	0.22 (0.19)	0.27 (0.19)
-3	0.27 (0.42)	0.3 (0.19)	0.27 (0.09)	0.14 (0.2)	0.23 (0.25)	0.24 (0.23)

3. Discussion

We summarise the key findings of the chapter, full replications of which are presented in appendix 2.2.

Summary of Findings

Recall versus Recognition for Overall Performance: It is striking that recall performance at 117ms is so low; see figure 3(B): the -1 item is correctly recalled on approximately one in five streams and, on average, a -1 to -3 item is recalled correctly on less than one in ten streams.

Additionally, although only an informal comparison, since they are based upon different measures – d-prime versus accuracy - it seems that recognition performance is substantially higher than recall, particularly at the 117ms SOA; compare figure 3 panels (A) and (B).

Recency and Regime Change: importantly, recall exhibits a substantially larger normalised recency effect than recognition at the RSVP SOA of 117ms; see figure 4. Comparing across the three SOAs, it seems that there is a large change in performance from 117ms SOA to 230ms SOA for recall (which might even be considered a regime change) that is not present for recognition. This is apparent in figure 3, but perhaps most dramatic in figure 4, where, for Recall, normalised recency is substantially larger for 117ms, with perhaps a small difference between 230ms and 350ms. In contrast, there is little evidence for a difference in normalised recency between the three SOAs for recognition.

Detection versus Recognition: recognition performance is clearly below detection; see figure 3(A). This difference is not as large as in Potter(1976), suggesting that our approach of probing memory for the most recently presented items in the stream has, as intended, provided a more conservative assessment of the difference between recognition and detection. Nonetheless, the difference is still substantial, for example, at 117ms SOA rates in the -1 position, recognition is 2.5 times detection.

Additionally, the difference between recognition and detection reduces from 117ms SOA to 350ms SOA; see figure 3(A). This suggests that, in relation to the memory trace left, the brain is particularly good at detection at fast SOAs, or in other words, detection and recognition memory tend to come more into alignment as the involvement of conscious perception increases.

Evidence in the literature for rapid memory formation: The low d’ in recognition and especially the extremely low accuracy in the free-recall task at the fastest presentation rate contrast with work claiming very rapid consolidation of working memory representations. Nieuwenstein and Potter (2006), showed that in whole-report--where participants report all items presented in RSVP--performance was relatively high (compared to a condition where only a subset of items had to be reported). Similarly, Vogel et al. (2006) , estimated from two experiments a rate of 50ms per item for visual working memory consolidation. A relevant difference between these studies and the present experiment is that we used a dual-task paradigm, where participants were instructed to search for a pre-specified target, while attempting to encode distractors for the memory tasks. This dual-task setting adds a degree of selection that may have impacted the consolidation of items resulting in more fragile memory traces. However, this dual task aspect is consistent with our interest in the memory traces left when performing detection.

Fragility of memory: A first indication of fragility is that, in general, memory traces at RSVP rates (117ms SOA) were weak: for the -1 position, which had the highest performance, recognition had a d-prime of 0.89 and recall had an accuracy of 0.23. However, memory fragility was most evident in the recency levels observed, with the substantially higher recency for recall, suggesting much greater fragility for recall than for recognition. This differentiation between recall and recognition is also supported by the big jump in performance from 117ms to the slower SOAs that is present for recall, but not so obviously for recognition. As we have discussed, this jump in performance may suggest a regime change for recall that is not present for recognition.

This differential fragility is likely to implicate working memory. At its heart, working memory offers representations that are freely accessible (Cowan, 2016). Thus, it may be that the fragility of recall arises, since, at RSVP rates, the stream is just going too fast to enable stable working memory representations to form. This difference in memory stability between recall and recognition resonates with the distinction in the long-term memory literature between recollection and familiarity (Wixted, 2007; Yonelinas, 2002). Recollection is typically associated with greater source memory, i.e. information about when or where the memory was formed. In contrast, familiarity may only involve a sense of knowing, without detailed contextual information about the eliciting experience. This also resonates with the Glance-Look model of the attentional blink (Barnard & Bowman, 2004; Su et al., 2011), which would suggest that many stimuli in RSVP are "glimpsed", but no more. Such glimpses may leave a sufficient trace for a sense of familiarity, but not for recall.

WM Capacity: Importantly, the poor memory retrieval that we observe is very unlikely to be impacted by working memory capacities, which would be expected to be considerably higher than we are observing. For example, if we add recall performances across the -1 to -3 positions, we only reach 0.8 of an item for SOA 350ms and 0.3 of an item for SOA 117ms. All proposed working memory capacities would be considerably larger than these estimates (Cowan, 2016).

Subliminal Salience Search (SSS): As previously discussed, our findings suggest a substantial disparity between recognition and detection performance. For example, detection was approximately 2.5 times recognition performance for the -1 item at a 117ms SOA and, as SOA increased to 350ms, this disparity reduced.

One reason for being interested in showing a difference between recognition and detection at RSVP rates (i.e. SOA 117ms) is that it suggests that the brain is able to find targets in RSVP, without leaving strong memory traces for the stimuli rejected as not targets. This, in turn, suggests that the brain is searching RSVP streams subliminally, as implied by the term Subliminal Salience Search (SSS), coined in Bowman et al. (2013).

The findings in Potter (1976) could be considered suggestive of subliminal salience search, but they left open the possibility of forgetting, i.e. that the vast majority of items presented in RSVP were consciously perceived, but were almost all forgotten before memory was probed at stream end. We have reduced the room for a forgetting explanation by specifically probing memory for items presented right at the end of the stream. For example, at the 117ms SOA, for the -1 position, the time from stimulus onset to recognition memory probe onset was 230ms. This leaves little time for the -1 item to be consciously perceived and then forgotten before memory is probed. Additionally, by randomly varying the length of streams, we prevented participants from being able to develop an expectation for when the distractor they would be probed on would arise in the stream.

STST and the Tokenized Percept Hypothesis: in the introduction to this chapter, we identified three hypotheses that arise from interpreting the Simultaneous Type/ Serial Token (STST) model as a theory of conscious perception. These concern the memory traces left by task-irrelevant distractors (i.e. non-targets) presented in RSVP. Our findings have provided evidence for all these three hypotheses: memory traces for irrelevant distractors 1) decayed rapidly; 2) had recognition performance substantially better than recall performance; and 3) had recognition recency stronger than recall recency.

Furthermore, the findings here could be argued to bring additional evidence for a theory that emerges from the STST model, which we have named the tokenized percept hypothesis (Chapters 1 and 2), and which is related to the linking of event individuation to awareness by Kanwisher (2001). In the Simultaneous Type/ Serial Token (STST) model (Bowman & Wyble, 2007) episodic information is associated with the occurrence of events by binding them to markers, called tokens. For example, these markers enable the system to retrieve the order in which events occurred and whether an event has been repeated, i.e. they encode classic episodic information.

Critically, in order to perform episodic tasks, such as identifying the order in which events occurred or whether they have repeated, sustained stable memory representations need to be maintained, e.g. the association between a token and an event. This is easily seen when considering how the brain would have to perform a pure repetition task, i.e. determine that a particular stimulus is repeating, when not told the stimulus that will repeat a priori. Critically, the brain only knows a stimulus has repeated at the point of that repetition, but to do this, it has to have successfully stored a memory trace of the stimulus' previous presentation. This requires durability of representations, something that the token system of STST provides.

Chapter 2 presented data suggesting that if not consciously perceived, representations of repeating items do not accumulate, i.e. they dissipate rapidly. Additionally, Bowman et al. (2014) provided evidence that when participants are looking for a target item in RSVP as well as for a repeating item, there is little, if any, evidence that the repeating item generates a P3 ERP component, even after 50 repetitions. In contrast, the target and an incidentally salient stimulus generated clear P3s. This suggests that participants were unable to create the durable memories for incidental items required to perceive repetitions.

The findings of this paper are wholly consistent with those of Chapter 2 and Bowman et al. (2014), all suggesting a marked fragility of representational traces elicited by fleetingly presented stimuli that are not identified as salient by the brain's attentional system, i.e. are non-targets. This raises the possibility that the fragility of pre-conscious memory representations prevents subliminal search on the basis of episodic properties, e.g. "find the repeating item". In turn, this raises the possibility that such episodic search can only be performed supraliminally, making it a capacity requiring the conscious brain. This is the tokenized percept hypothesis.

In this sense, one could argue that subliminal salience search suggests one thing that the subconscious is good at: searching for stimuli that are known and a priori salient, while the tokenized percept hypothesis suggests one thing it is not good at: searching for stimuli purely on the basis of the episodic properties of their occurrence in the world.

Chapter 4 delves extensively into the phenomenon of non-reportability of items in Rapid Serial Visual Presentation (RSVP) tasks. To mitigate potential confounding factors, two specific measures were implemented. Firstly, the retention interval between the critical item and the probe was minimized by evaluating the last item (preceding the mask) within the sequence. Secondly, the reportability of task-irrelevant items in RSVP was assessed through the establishment of conditions conducive to incidental encoding.

Chapter 4: The Missing Memory Imprint

1. Introduction

Bowman et al (2013) introduced the term *subliminal salience search* (SSS) to describe the capacity that we have to detect and identify salient stimuli in our environments, with perception of salient stimuli in Rapid Serial Visual Presentation (RSVP) streams used as a representative case. Clearly, the loaded term here is *subliminal*: do we definitively know that the brain is finding salient stimuli in RSVP streams below the threshold of awareness? Indeed, it could be that the term *liminal* salience search is more appropriate, i.e. that percepts are fleetingly experienced in RSVP, but are not strictly subconscious. The objective of this paper is to provide evidence that the sub- in subliminal can be appropriately used.

The perceptual experience associated with viewing RSVP is certainly strange, with a jumble of stimuli appearing rapidly, with vividness of percept varying according to changes in masking strength between items. However, objectively measured performance at seeing salient stimuli, across a range of possible dimensions of salience, is excellent. A classic example of such high performance is Potter et al, 2014, see figure 1, where participants searched RSVP streams of pictures on the basis of meaning with stimuli presented at 78 per second; other examples include (Potter, 1976; Barnard et al, 2004; Most et al, 2005; Bowman et al, 2013), and many others.

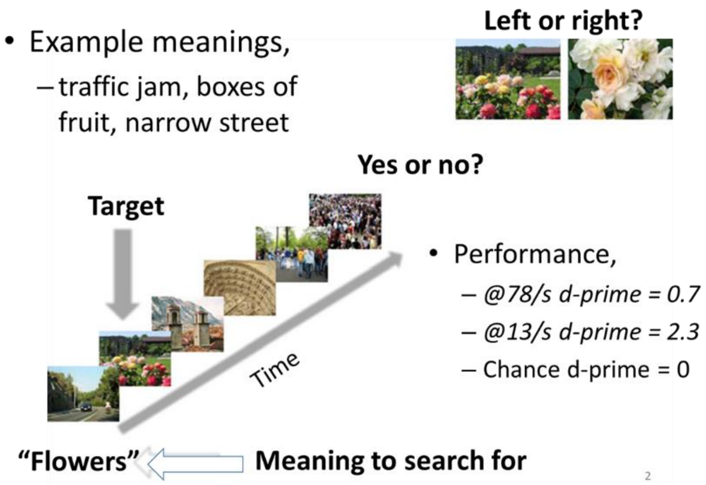

Figure 1: Illustration of brain's great capacity to search RSVP streams. Participants were given a meaning to search for in an RSVP stream of pictures. At stream-end, they performed a presence/ absence judgement on the occurrence of a picture consistent with the target meaning and then a 2-alternative forced choice on its identity. Notably, even when pictures were presented at 78 per second, detection performance was well above chance, with a d-prime of 0.7. Figure reprinted from Potter, M. C., Wyble, B., Hagmann, C. E., & McCourt, E. S. (2014). Detecting meaning in RSVP at 13 ms per picture. Attention, Perception, & Psychophysics, 76(2), 270-279.

Thus, it is clear that the human brain is very good at finding salient (e.g. target) stimuli in RSVP streams. In this sense, the brain is performing a search when viewing such streams, i.e. it is seeking out stimuli that are of interest to it. When such a stimulus is found, it is more fully perceived and then enters into working memory, enabling it to be reported.

The key logic underlying the approach taken in this paper is that if the process of finding salient stimuli in RSVP is conscious, the case of liminal salience search, memory imprints should be left for many of the not salient stimuli presented in the stream. In terms of classic RSVP experiments, this suggests interest in the memory traces left by arbitrary distractors.

In particular, we can follow the logic that we introduced in Chapter 3, which proceeds as follows. Since the position at which targets can appear in streams is randomly varied across trials, a high detection performance, where, say, 90% of trials are correctly responded to (hits or correct rejects), would require a very high percentage of the items presented in the stream to be processed sufficiently to determine if they are targets. If they are judged to be targets, they are encoded into working memory, but what happens to stimuli that are not judged to be targets? Under the liminal salience search hypothesis, some sort of memory trace would be left for a very large number of items, even if only one of these was a target.

The question of the memory trace left in RSVP was first investigated in Mary Potter's classic work in the 70's (Potter, 1976). More recently, we showed that free recall was very poor for RSVP stream distractors (Chapter 3). However, neither of these previous studies probed memory in a fashion that could truly extrapolate to an arbitrary distractor in an RSVP stream, which is what is of interest for the SSS hypothesis. Thus, the central issue remains moot. This is what we seek to resolve in this article.

In the studies reported here, we probe recognition and detection, two yes/no decision tasks, enabling us to compare d-primes for detection to those for recognition. However, this does not resolve the key experimental difficulty, which is that the memory traces left could be very brief, i.e. the percept created by a distractor could be forgotten very rapidly. That is, strictly speaking, there was an imprint generated, but it evaporated long before being probed for at the end of the RSVP stream.

We resolve this issue by, firstly, randomly varying the length of RSVP streams, in order that participants cannot build up an expectation for when the end of stream would come, and, secondly, by just probing memory for the last distractor presented. This gives the shortest time for forgetting before probing memory. Although, of course, this distractor has to be masked.

We explore two Stimulus Onset Asynchronies (SOAs), 117ms and 350ms. The first of these is a typical RSVP rate, with stimuli presented “on the fringe of awareness”, or, indeed, below it, while the second should be a conscious regime, in which the vast majority of stimuli would generate a conscious experience.

Our first experiment compares recognition to detection for a range of (last item) masks – words, pseudo-words and random symbol strings – as well as a “boundary” condition in which no mask was presented. This enables us to determine whether reduced recognition performance compared to detection was robust across types of mask, and additionally whether masking was more “perceptually-determined” (as opposed to conceptually) at short, rather than at long SOAs.

Our second experiment seeks to test for memory imprints, when the brain is in its “native” (sub)liminal search state. Specifically, we are interested in the memory traces left by distractors, when the brain is solely engaged in search for target stimuli, i.e. is just performing the detection task. Whether instructed to do so or not, as soon as memory has been probed once during an experiment, participants will be attempting to encode distractors and they will effectively be performing a dual task: “detect targets *and* encode distractors”. To resolve this issue, we use the approach introduced in Chen and Wyble (2015), and introduce a single surprise recognition memory probe during a pure detection experiment.

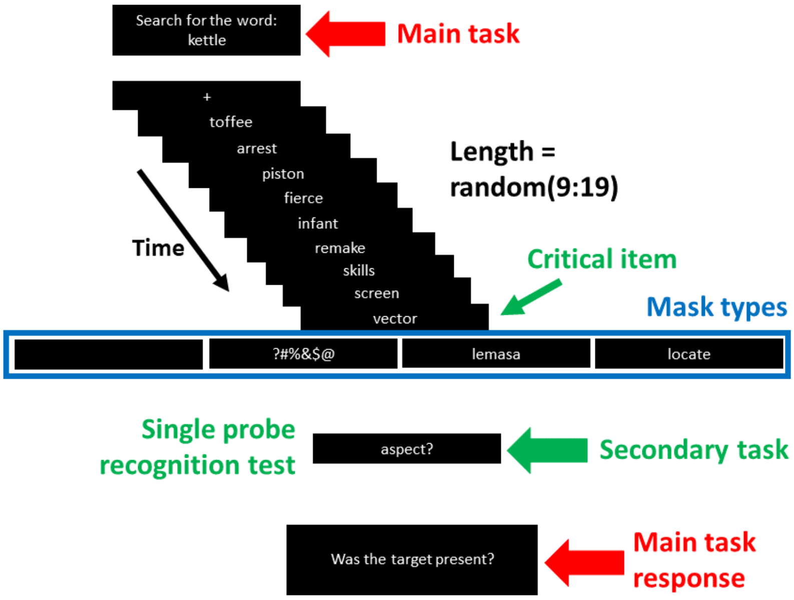

Figure 2: experiment 1 procedure: RSVP streams of words were presented, terminated by a mask. Streams were stopped randomly from the 9th to the 19th item. The main task was to detect the specified word and perform a presence/ absence judgement on it at stream end (see Main task response). Additionally, in a Secondary task, participants performed a single-probe recognition test on the critical item: the last stream item before the mask. Four different mask types were used.

2. Experiment 1

Methods

Participants

30 (mean age=19.33, sd=1.24) undergraduate students (8 males) of the University of Birmingham took part in the experiment in exchange for course credits. All were native English speakers and had normal or corrected-to-normal vision. The experiment

conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Materials

1000 English words were selected from (Warriner, Kuperman, & Brysbaert, 2013) to serve as distractors. 224 from the same source were selected as Targets for the detection task. 224 words were selected for the recognition test (half served as Old Words and half as New Words). Four lists of stimuli were created so that the Old Words were presented in each of the masking conditions. Additionally, we generated 56 pseudo words by changing two letters of each of a new set of 56 words.

Procedure

We presented RSVP streams on a 24'' LCD screen (refresh rate: 60Hz, resolution: 1920 x 1080) using custom Psychopy3 (Peirce et al., 2019) scripts running under Python 3.6. Stimuli were Arial white characters of 1.3° of visual angle in height on a grey background (40% white). Participants were seated at 60cm from the screen. The experiment was conducted individually in a quiet room. Each experimental session consisted of 224 RSVP trials.

The format of the experiment is shown in figure 2. Each trial started with the presentation of the instruction "Search for the word:" in the upper part of the screen along with the target word presented at the centre of the screen. Participants were instructed to press the "space bar" after reading the target word to continue. The starting item was a fixation cross (+) presented for 500ms at the centre of the screen. Then a sequence of 9 to 19 (varied randomly) words was presented at an SOA of 117ms or 350ms. The finishing item, presented for the same SOA as the other stream items, varied depending on the masking condition:

Word: a random word from the main list of distractors

Pseudoword: one of the pseudowords described in the Materials section.

Symbols: a combination of these symbols &#?@$% in random order

Blank: a blank screen (no stimulus) for the same SOA duration.

The response phase began immediately after the presentation of the finishing stimulus. Participants answered two successive questions. First, the Recognition question "Have you seen this word?" was presented in the upper part of the screen along with a probe word presented at the centre of the screen. The probe consisted of the word presented immediately before the finishing item (the mask) in half of the trials and was a new word in the other half. Participants pressed with the right hand the "left arrow key" to indicate YES or the "right arrow key" to indicate NO. Second, the Detection question "Have you seen the target?" was presented and they responded YES or NO using the same keys they used for the Recognition question.

Results

The first objective of this experiment, as stated above, was to demonstrate that detection performance was disproportionally better than recognition for the short SOA (117ms) relative to the long SOA condition (350ms). To do that we obtained the d-prime values from the Hit and False Alarm rate for both tasks: detection and recognition, see Figure 3. A three-way (Task: detection and recognition, SOA: 117 and 350, Mask: blank, symbols, pseudoword and word) within-subjects ANOVA showed that Task interacted significantly with SOA $F(1,29)=10.46$, $MSE = 4.96$, $p<0.01$, $\eta^2=0.27$. Pairwise comparisons showed no effect of SOA in the Detection task ($p>0.1$), and a significant effect of SOA in the Recognition task, $t(1,29)=7.4$, $p<0.001$, $d=1.2$. Task also interacted significantly with Mask, $F(3,87)=14.29$, $MSE = 3.68$, $p<0.001$, $\eta^2=0.33$. In addition, we found that the three main effects reached significance: SOA, $F(1,29)=50.26$, $MSE=14.34$, $p<0.001$, $\eta^2=0.63$; Task, $F(1,29)=106.33$, $MSE=61.64$, $p<0.001$, $\eta^2=0.79$; and Mask, $F(1,29)=23.16$, $MSE=6.83$, $p<0.001$, $\eta^2=0.44$. Finally, the interaction between SOA and Mask approached significance $F(3,87)=2.43$, $MSE=0.64$, $p=0.07$, $\eta^2=0.08$. The three way interaction between Task, SOA and Mask was not significant $F(3,87)<1$, $p>0.1$.

The second, more exploratory, objective was to examine to what extend the last stimulus presented in the streams modulated recognition performance in both SOAs. Additionally, to design Experiment 2, we aimed to identify which type of mask would result in higher recognition performance, while still showing the basic interaction between Task and SOA. In this respect, the blank condition (unmasked critical item) was of no interest, since recognition performance was sufficiently high that the interaction between SOA and Task was not significant ($p>0.1$). Additionally, unmasking the critical item would have made it unrepresentative of the masking of arbitrary distractors in RSVP. For the other masking conditions, the planned pairwise comparisons on the d-prime values of the Recognition task showed that for the 117ms SOA there were no significant differences in all the comparisons among the Words, Symbols and Pseudowords conditions (all $ps>0.1$). In contrast, in the 350ms SOA condition, the Symbols mask showed larger d-prime values relative to the Words $t(29)=2.32$, $p=0.03$, $d=0.42$ and the comparison between Symbols and Pseudowords approached significance $t(29)=1.92$, $p=0.06$, $d=0.35$.

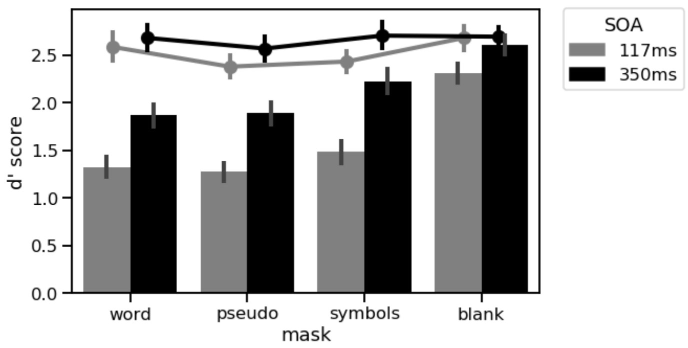

Figure 3: results of experiment 1, comparing recognition (bars) to detection (dots) performance across masking conditions and SOA.

As previously discussed, experiment 1 compares detection and recognition performance in a dual task context, i.e. in which participants will be seeking to trade performance on the two tasks off against each other. However, a key question is the memory trace left for an arbitrary distractor in RSVP when participants are solely engaged in detection. Experiment 2 considers this question by incorporating a surprise memory probe, see figure 4.

On the basis of the results of Experiment 1, we selected the symbols mask. The logic here was to use the masking that gives the most capacity to observe recognition memory for distractors. Thus, we selected the mask with the highest memory performance in Experiment 1. If we are able to exhibit low memory performance even for the weakest mask, we will have given the greatest capacity for a memory trace to be observed if it is there.

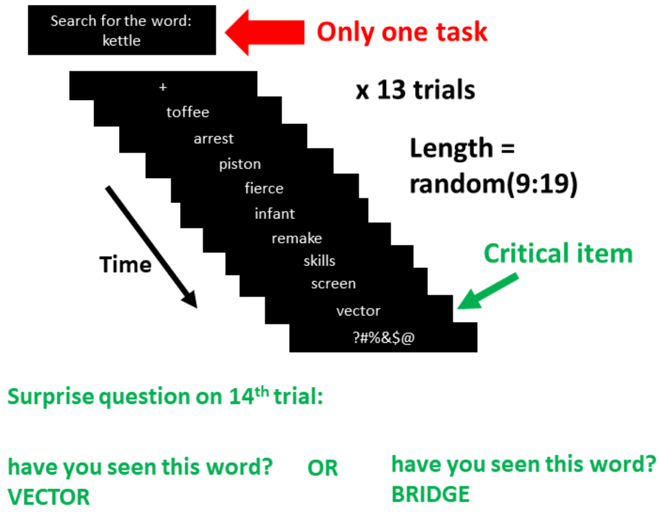

Figure 4: experiment 2 procedure: RSVP streams of words were presented, terminated by a mask. Streams were stopped randomly from the 9th to the 19th item. The task was to detect the specified word and perform a presence/ absence judgement on it at stream end (not shown). Additionally, a surprise question arose on the 14th trial, asking participants to perform a single-probe recognition test on the critical item: the last stream item before the mask.

3. Experiment 2

Methods

Participants

151 undergraduate students of the University of Birmingham took part in the experiment in exchange for course credits. 60 (mean age= 19.23, sd= 1) participants (9 males) took part in the lab session of the experiments. The rest, 91 (mean age= 19.65, sd=

1.48), took part in the online session. All were native English speakers and had normal or corrected-to-normal vision. The experiment conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Materials

Distractors were selected randomly from the main list of distractors of Experiment1. In addition 30 words were selected for the Recognition test (a list of 15 items to serve as Old words and 15 as New Words). Finally, target words for the Detection task were selected from a list of 30 words.

Procedure

For the lab version of the experiment, screen model, stimuli presentation software, and stimuli presentation parameters (font and background) were the same as in Experiment 1. The online version was created using PsychoPy (Pierce et al 2019) and hosted by https://pavlovia.org/. The procedure is shown in figure 4. Participants were informed that the experiment consisted of a series of trials, where they would be presented with rapidly presented streams of words. Their task was to search for a specific target word that could be inserted in the streams. At the end of the trial, they were expected to indicate whether they have seen the target.

The experiment consisted of 14 RSVP trials. The structure of each trial was the same as in Experiment 1: Instruction “Search for the word: <target>”, fixation cross, stream of words, (symbols) mask. The SOA of the streams of words was 117ms. The crucial manipulation of this experiment was in the response phase. The first 13 trials ended with the presentation of a Detection question “Have you seen the target word”. As in Experiment 1, participants had to press the “left arrow key” for YES and the “right arrow key" for NO. Crucially, the last trial, 14^{th}, was followed by a surprise Recognition question. In this trial, immediately after stimuli presentation, participants were presented with the text "SURPRISE QUESTION, Have you seen this word?" presented at the upper part of the screen along with a probe word presented at the centre of the screen, and the text "press <- to indicate YES or -> to indicate No" presented in the lower part of the screen. The experiment ended after the participants responded to this question.

Data analysis

Data consisted of the subjects' responses to the surprise recognition question that ended Experiment 2. 2 subjects were excluded for taking too long to respond (more than 5 seconds). Out of the 58 subjects (and data points) 23 had been presented with a probe corresponding to a New Word and 35 had been presented an Old Word. In addition, out of the 91 subjects who did the online experiment, 42 had been presented with a probe corresponding to a New Word and 49 were presented with an Old Word. We used a model-based approach to estimate Signal-Detection Theory parameters (DeCarlo, 1998), avoiding the need to collapse the data (to get proportions of Hits and False Alarms to obtain the d-prime). In consequence, the responses of the participants were fitted with a generalized linear model (with a Probit link function). The data analysis strategy was designed to test to what extent the data supported the Null Hypothesis of a d-prime value of 0. To do this, as shown in figure 5, we built a Bayesian generalized linear regression model to estimate a posterior distribution of d-prime values. Then, the Bayes Factor was obtained using the Savage-Dickey density ratio (Wetzels et al., 2009) between the prior and posterior estimates.

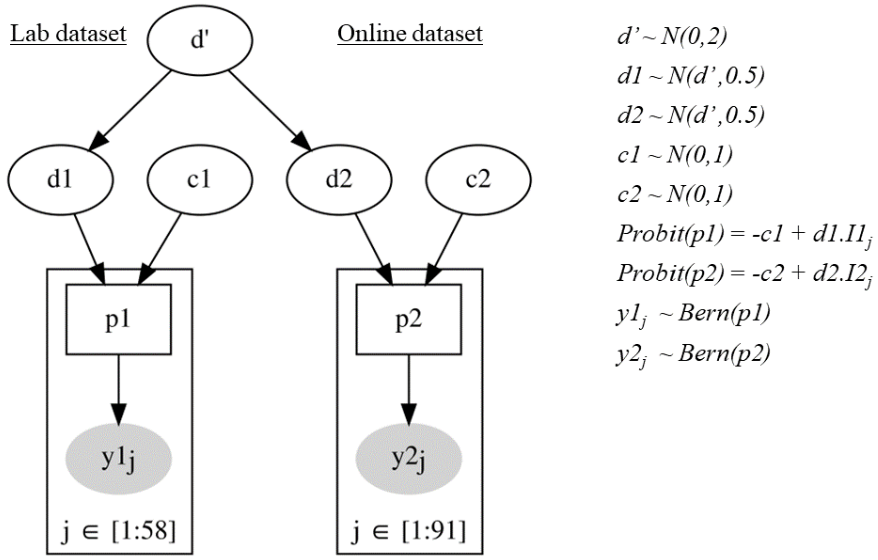

Figure 5. Specification of the Bayesian model for Experiment 2. Index 1 refers to the in-person (Lab) experiment and index 2 the online version. For each participant (j), $y1_j$ and $y2_j$ are either a one or a zero, depending upon whether the jth participant responded with YES or NO on their 14th trial. p1 and p2 are probabilities of a YES response across all 14th trials. d′ is the highest level d-prime, which is a prior for d1 and d2, the d-primes for each of the two datasets (in-person and online). c1 and c2 are the criteria for the two datasets, and I1j and I2j indicate target presence/absence on 14th trial of participant j. The negative of each criterion gives the intercept, and the relevant d-prime, the slope of the probit regression. *Bern(pi)* indicates a Bernoulli distribution, for which 1 has a probability of pi and 0 a probability of (1-pi).

The second objective of the analysis was to test whether the results of Experiment 2 were significantly different from those of Experiment 1. To do this, we selected the data from Experiment 1 where the same SOA and Mask conditions as in Experiment 2 were used. To fit the data from Experiments 1 and 2, we extended the Bayesian model from

Experiment 2. The part of the model that fits the data from Experiment 1 is a Hierarchical version of the model for Experiment 2, where each participants’ parameters are drawn from parent distributions. d-prime values were estimated for both experiments. The Bayes Factor was then obtained by using the Savage-Dickey density ratio between the Effect Size density estimated from the priors and from the posterior at 0.

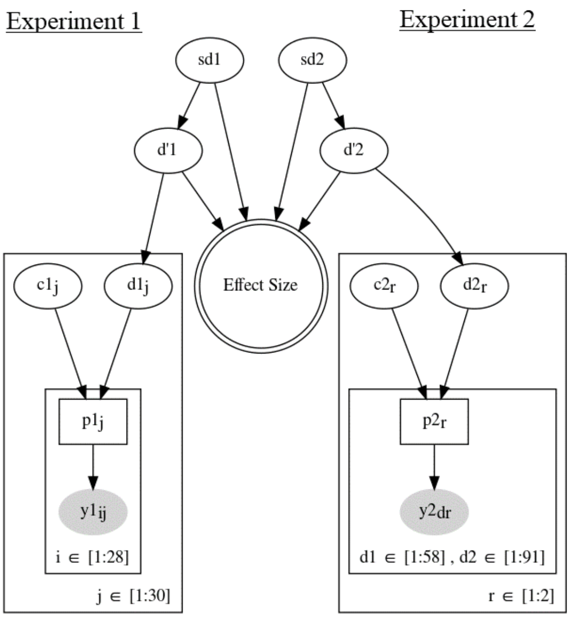

Experiment 1

$sd1 \sim HalfStudentT(\sigma=2, nu=10)$

$d'1 \sim N(0, sd1)$

$c1_j \sim N(0,1)$

$d1_j \sim N(d'1, 0.5)$

$Probit(p1_j) = -c1_j + d1_j . I1_j$

$y1_{ij} \sim Bern(p1_j)$

Experiment 2

$sd2 \sim HalfStudentT(\sigma=2, nu=10)$

$d'2 \sim N(0, sd2)$

$d2_r \sim N(d'2, 0.5)$

$c2_r \sim N(0,1)$

$Probit(p2_r) = -c2_r + d2_r . I2_r$

$y2_{dr} \sim Bern(p2_r)$

Effect Size $= d'1 - d'2 / \sqrt{sd12 + sd22}$

Figure 6. Bayesian hierarchical model for Experiments 1 and 2. Index 1 refers to the Experiment 1 and index 2 Experiment 2. In Experiment 1, for each participant (j) and trial (i), $y1_{ij}$ is either a one or a zero, depending upon whether the participant responded with YES or NO on the recognition question. In Experiment 2, for each version of the experiment (r, in-person and online) and each participant (d, which corresponds to d1 in first version of experiment and d2 in second version), $y2_{dr}$ is either a one or a zero, depending upon whether the participant responded with YES or NO on the 14th trial. $d'1$ and $d'2$ are the highest level d-prime from both experiments, which are priors for $d1_j$ and

$d2_r$, the d-primes for each participant in Experiment 1 and for each dataset (in-person and online) in Experiment 2. $c1_j$ is the criterion for each participant of Experiment 1, and $c2_r$ is the criterion for each dataset of Experiment 2. $I1_j$ and $I2_r$ indicate target presence/absence. The negative of each criterion gives the intercept, and the relevant d-prime, the slope of the probit regression. *Bern*(p_j) and *Bern*(p_r) indicates a Bernoulli distribution, for which 1 has a probability of p_j (or p_r, for *Bern*(p_r)) and 0 a probability of (1-p_j) or (1-p_r, for *Bern*(p_r)).

Results

Evidence for the Null Hypothesis (d-prime=0)

For the lab version the proportion of Hits was 0.57 and the proportion of False Alarms was 0.43 (d-prime: 0.34). For the online sample, the proportion of Hits was 0.61 and the proportion of False Alarms was 0.45 (d-prime = 0.40). To estimate the posterior of the parameters, 5000 draws (tuning samples = 300) were sampled from it using the No U-Turn Sampler (NUTS) (Hoffman & Gelman, 2014). With a Normal distribution (mean=0, sd=2) as a prior for the d-prime, a Bayes Factor of 3.29 (see Figure 7) was obtained in favour of the Null Hypothesis and the mean of the posterior of the d-prime was 0.36 (94% credible interval = -0.41 to 1.12). In addition, to assess the robustness of the Bayes Factor, we fitted the model using a Uniform prior (lower bound=-3, upper bound=3). In this case, the Bayes Factor in favour of the Null Hypothesis (d-prime=0) was 3.84. The mean of the posterior of the d-prime was 0.38 (94% credible interval = -0.38 to 1.17).

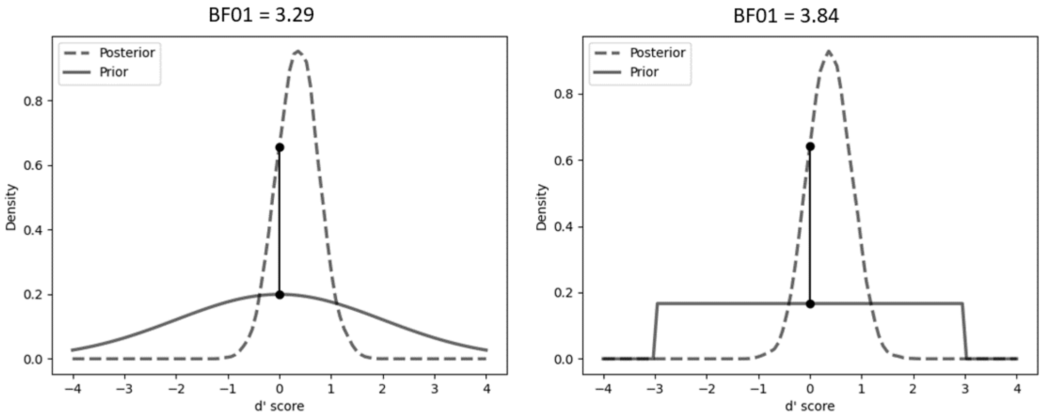

Figure 7. Posterior (dashed line), prior (solid) and Bayes Factor (density ratio, at vertical line at 0), Experiment 2. Left, Normal Prior (mean=0, std=2). Right, Uniform Prior (-3,+3).

Difference between Experiments 1 and 2

The standardized effect size of the difference between the d-prime estimates of Experiments 1 and 2 was obtained by sampling 3000 draws (tuning=300 draws) from the posterior of the model depicted in Figure 6. Similarly, 3000 draws were obtained to estimate the Effect Size from the priors of the model. The Bayes Factor was the ratio at 0 (no difference between Experiments) between the estimates of the Effect Size from the posterior and priors. The Bayes Factor was 12.87 (see Figure 8), which indicates strong evidence for the Alternative Hypothesis of a significant difference between the performances of the two experiments.

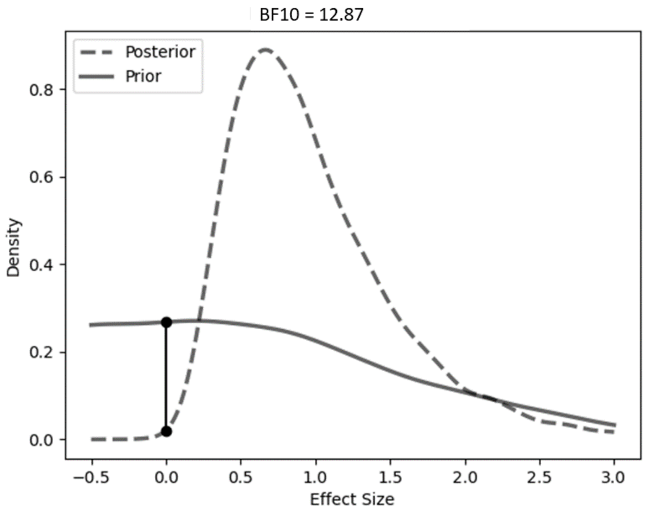

Figure 8. Posterior distribution (dashed line), prior distribution (solid) and Bayes Factor (density ratio, at vertical line at 0) of the Effect Size of the comparison between Experiment 1 and 2.

4. Discussion

Summary of findings

In these two experiments, we compared detection (for targets) and recognition (of distractors) performance in RSVP. An important finding was an interaction between Task and SOA in Experiment 1, whereby the change in SOA did not impact Detection, but substantially impacted Recognition. This replicates the finding in Chapter 3, where it was also argued that the interaction is unlikely to be driven by a ceiling effect on Detection performance. This interaction sits well with our perspective that Detection performance is high when the brain is searching subliminally for targets (SOA 117) and not much higher, if at all, when searching supraliminally (SOA 350), indeed discrimination (d') does not change between these two SOAs. However, recognition improves dramatically from SOA

117 to SOA 350, suggesting a regime change for recognition between subliminal and supraliminal. This fits with our theoretical perspective that the brain can search subliminally for what it "knows", this is subliminal salience search, but consciousness is important when laying down long lasting memory traces.

It is also the case that as expected, d-prime values were significantly lower for Recognition compared to Detection (see Figure 3) independently of the SOA or the type of mask that followed the critical items. Furthermore, recognition performance was clearly above chance in Experiment 1 even for the shortest SOA. However, in Experiment 2, we found evidence that recognition performance is extremely low (see Figure 7, for the posterior estimates of the model) when the recognition question was not expected by the participants. That is, in the context of a single task (Detection), recognition performance was on the basis of a Bayesian test, not different to chance, and significantly lower than the d-prime values obtained when Recognition was part of the task-set, as in Experiment 1.

Comparing recognition and detection

Central to this paper is our comparison of d-primes for detection and for recognition. How can we conceptually justify this comparison, since these are two different cognitive functions? Essentially, d-prime in RSVP, indexes the capacity to discriminate the presence of a critical item appearing at a single arbitrary position in a stream (quantified as the Hit rate) versus not appearing at any position in the stream (quantified as a Correct Reject). The difference between the two is that, for detection, the critical item is the target being searched for, while, for recognition, it is the (single) item probed in the recognition memory display. In addition, although our recognition test probes the last item (before the mask) in the RSVP stream, our approach of randomly stopping the stream, to prevent any capacity to predict when the end is coming, simulates (in the mind of the participant) the occurrence of the probed item at any arbitrary position in the stream.

Thus, essentially, the detection and recognition d-primes are assessing the same discrimination, with the only difference being whether the critical item is the (detection) target or the item "post probed”.

Relationship to our Previous Work

Chapter 3 was the first to extended Potter (1976)’s findings by firstly, probing memory specifically for items at the end of the stream, in order to reduce the possibility for forgetting, and secondly, by varying the length of RSVP streams, in order to prevent participants from predicting their ending. Chapter 3 focussed on characterising recency and comparing it between recall and recognition. A striking finding was that at RSVP rates, recall was, firstly, extremely poor and secondly, exhibited a close to maximal recency effect, in the sense that the little that was recalled was almost exclusively at position -1, i.e. almost no items at positions -2 and -3 were recalled.

Most importantly, though, the poor recall performance identified in Chapter 3, suggests that the brain’s capacity to detect and identify targets in RSVP, with detection performance often around 80%, cannot be justified by a recallable level of memory imprint. In Chapter 3, the -1 item could be recalled on less than one in five streams (i.e. <20%).

This excludes free recall as a source of a memory imprint left by distractors during the process of determining that they are not targets. However, recognition performance was higher in Chapter 3, leaving the possibility that rejected distractors leave a familiarity, rather than recollection trace, an idea that would be consistent with the glance phase of the Glance-Look model of RSVP perception Su LI et al (2011). This possibility that there is a familiarity trace is what we have specifically focused on in this paper.

Additionally, we refined the experimental paradigm employed in Chapter 3. Firstly, we only probed recognition memory, recall was not considered. Secondly, we probed memory with a single item at the end of each RSVP stream. This is to avoid the disruption that can arise in multiple-item probes, where each memory retrieval can potentially disrupt the memory for other items being probed; see, for example, the effect of retro cues in iconic memory experiments (Sligte et al, 2008; Sperling, 1960). Thirdly, we only probed the -1 item (i.e. the last before the mask) in the RSVP stream, as a result of which we have not considered recency in this paper. Finally, we explored different types of end of stream mask.

In pilot work, we determined that the mask used in Chapter 3, a sequence of hashes, gave substantially higher performance than variable-character masks. However, the central idea of our experiments is that the distractor that we probe should be presented with a similar perceptual difficulty to an arbitrary distractor in the stream, which would be backward masked by another distractor. Accordingly, here we preferred an end of stream item that is a variable-character mask. We explored different possible letter masks in experiment 1 and selected the symbols mask, since it has the highest performance across the variable-character masks, enabling us to give as much chance of high performance coming out as possible.

Surprise Experiment

Our second experiment seeks to assess the memory left for an arbitrary RSVP item, while the brain is in its “native” subliminal search state, i.e. it is just performing a detection task. Although we have presented evidence for zero memory for RSVP distractors (Bayes Factor in favour of d-prime=0 was 3.84), we do acknowledge the need for a replication of this effect. Indeed, it seems surprising that there really is a complete absence of a memory imprint. However, what we believe we can argue is that the memory imprint is at least weak, and is certainly very substantially smaller than detection performance at RSVP rates. This is sufficient for the line of argument we are making here, i.e. that the brain can exclude distractors in RSVP with very little in the way of a memory trace being left for those distractors.

An alternative explanation of the poor performance in the surprise recognition question is that this is a simple case of forgetting. The retention interval (the gap between presentation and probe) and the presence of the surprise question might be sufficient to disrupt the retrieval of the critical item. More evidence is required to reject this hypothesis, but initial evidence seems to rest credibility from the forgetting explanation. Chen and Wyble (2016), using a similar surprise question method, showed successful retrieval of perceptual attributes that participants did not expect to report. That is, the surprise question did not disrupt the memories for attributes that needed to be encoded for other purposes than report. They concluded that the poor performance in a surprise recognition question demonstrates a failure of memory consolidation rather than forgetting.

Furthermore, we believe that the brevity of the retention interval might be sufficient to reject the forgetting hypothesis. Indeed, evidence from RSVP Event Related Potential (ERP) experiments suggest latencies for conscious perception. In RSVP, the ERP component most commonly associated with conscious perception is the P3b (Bowman et al, 2013, 2014; Craston et al, 2009; Vogel et al, 1998; Sergent et al, 2005); see, for example, Pincham et al (2016) and Jones et al (2020) for particularly direct support for this association. Across the numerous RSVP ERP experiments, the P3b is rarely if ever, found to start before 350ms and to peak before 450ms after the eliciting stimulus.

If we extrapolate these P3b latencies to our experiments, we are left with the scenario that conscious perception of the -1 item would occur perhaps 170ms after the memory probe started being presented[6]. In this context, is the forgetting hypothesis still tenable? It would require the -1 stimulus to be perceived and forgotten all while the memory probe screen was being displayed, and indeed had already been presented for some 170ms before the perception-forgetting would start unfolding.

Conclusions

The findings presented here and in chapters 2 and 3, suggest a marked fragility (indeed potentially absence) of memory for incidental stimuli presented in RSVP. It is clear that *salient* stimuli in RSVP streams frequently create durable traces, enabling accurate presence/absence or identification judgements on targets. However, on the basis of our work, it seems that the brain can propel stimuli of interest into consciousness and working memory, without incurring a memory storage "cost" for stimuli that are not of interest. In this sense, the brain successfully applies a strong salience filter (Lachter, Forster, & Ruthruff, 2004; Bowman & Wyble, 2007).

We have talked in terms of memory for incidental stimuli being a "cost". This is true if one views the task of the perceptual system as to act as a filter; however, the quotation marks here are pertinent, since this failure in incidental encoding does rule out one valuable functional capacity, and that is, the capacity to attribute episodic information to

[6] The latency of the P3b suggests an approximate time to conscious perception of 400ms from the -1 item onset. The memory probe onsets 230ms after the onset of the -1 item. 400 minus 230, gives us 170ms.

stimuli. As an illustration, let's take the ability to know that a stimulus has been previously presented (i.e. that it is a repetition) as a classic example of an episodic task. A repetition task can only be effectively performed if there is *incidental durability*. That is, the stimulus only becomes salient at the point of its repetition; however, to know that it has been repeated, a previous instance of occurrence has to have left a memory trace, and at the point of that first occurrence it was *not* salient. Does this indicate a key duality between the subconscious and the conscious: the subconscious can find stimuli in our environments with exquisite accuracy, while only the conscious can exhibit incidental durability enabling episodic information to be represented? The former of these is subliminal salience search, and the latter what we call the *tokenized percept hypothesis* Chapters 1 and 2, which has similarities to a hypothesis that Kanwisher has presented (Kanwisher, 2001). At its heart, the tokenized percept hypothesis claims that a central function of conscious experience is to episodically tag experiences in time and indeed space, and that capacity is uniquely conscious.

Chapter 5 adopts a Proactive Interference design to investigate the matter of episodic encoding of stimuli in Rapid Serial Visual Presentation (RSVP) tasks. Despite previous evidence indicating a low reportability of items in RSVP (as demonstrated in Chapters 3 and 4), it remains plausible that the memory representations for these items may be too feeble to be consciously reported, yet possess sufficient strength to influence subsequent items. Thus, Chapter 5 aims to address this inquiry by examining the phenomenon of the list length effect in RSVP in a Recognition task.

Chapter 5: Stimuli Presented on the Fringe of Awareness Evade Proactive Interference

1. Introduction

Can episodic memory operate outside consciousness? Episodic memory is a form of declarative memory and, as such, it is assumed that consciousness is required for the encoding of events and the retrieval of memories (Squire & Zola, 1996). In fact, consciousness has been used as a criterion to differentiate between memory types: declarative and nondeclarative. This strategy has been challenged by evidence showing the possibility of episodic memory formation for stimuli presented subliminally (Chong et al., 2014; Duss et al., 2014; Reber & Henke, 2011). This evidence fits well with the increasingly popular view postulating that consciousness is not required for many higher-order cognitive processes (Hassin, 2013), and episodic memory might not be an exception (Henke, 2010).

Investigating these issues often requires the use of experimental paradigms where stimuli are presented on the fringe of awareness. One of these paradigms is Rapid Serial Visual Presentation (RSVP). In RSVP, items are presented sequentially in the same position of the screen at a fast pace (around 10Hz). Using this paradigm, one can present images or verbal stimuli and use a memory test, such as a Recognition or Free Recall, to evaluate the strength of the memory traces laid down by the RSVP presentation. Since Potter (1976), we know that performance in a Recognition test is quite poor relative to the performance in identification and detection tasks. This failure of reportability is consistent with a view that many items in RSVP steams that are not searched-for may fail to be consciously perceived (Chapter 2).

Indeed, RSVP has been extensively used to investigate conscious perception, giving rise to a rich literature that has shed light on the interplay between attention and consciousness. The Attentional Blink and Repetition Blindness illustrate how, when perception is at the limit of its processing capacity, even task-relevant stimuli are missed and not reported. In the Attentional Blink, observers fail to perceive a second target (T2) when it is presented in close succession (200-500 ms) after a first target (T1) (Raymond et al., 1992). Similarly, Repetition Blindness is the failure to report the second occurrence of a repeated item in RSVP (Kanwisher, 1987). These failures of perception are often interpreted as requiring the conscious encoding of the T1, in the Attentional Blink, or the first occurrence of the repetition in Repetition Blindness (Kanwisher, 1987; Nieuwenstein et al., 2009).

However, the use of RSVP to study conscious perception is not exclusive to the Attentional Blink and Repetition Blindness experimental manipulations. In fact, based on the low reportability of regular items in RSVP (Chapter 2, Nieuwenstein & Potter, 2006; Potter, 1976), it can be argued that failures of perception are bound to happen for some items regardless of the identity or the task-relevance of prior items in the stream. Do these reportability failures indicate that regular items in RSVP streams might not be consciously perceived? There is not a straightforward answer to this question. One of the issues with paradigms such as RSVP is that reportability requires the conscious encoding of an item and the memory capacity to recall it. This memory capacity condition is not trivial in RSVP experiments, where typically the number of items presented clearly exceeds the span of short-term memory (Cowan, 2010). This limitation does not affect AB or RB since just two items need to be detected/identified and reported at the end of the stream, and performance is compared with the reports of targets presented in later positions outside the critical window. However, when all items are task-relevant, the "memory confound" remains, and, handling it, might require looking beyond raw comparisons in reportability performance.

As mentioned above, we are interested in testing the possibility of episodic memory operating outside consciousness. This hypothesis entails that below-threshold– non reportable – stimuli could successfully be encoded in memory. These memory traces, too weak to support reportability, might be strong enough to contribute to some of the classic effects in episodic memory, such us repetition (Ebbinghaus, 1964) or interference (Müller, & Pilzecker, 1900) effects. These effects can be studied for stimuli presented in RSVP. In fact, we have addressed the issue of conscious perception and repetition using RSVP (Chapter 2). Here we will explore the case of interference effects.

Perceived stimuli that are encoded in memory might still fail to be reported in a recall or recognition test due to the presence of competing information that comes to mind when cued. This interference effect is perhaps the most important factor responsible for forgetting (Kahana et al., 2022). That is, the participant's studying of additional information, along with the target information, impairs the retrieval of the critical items. If the additional information is presented *after* the target, there is Retroactive Interference (Müller, & Pilzecker, 1900). When the additional information is placed before the target then we observe Proactive Interference (PI) (Underwood, 1957). In Chapter 3, we have explored Retroactive Interference effects in RSVP by looking at the recency curves (how performance is modulated by the position of the last items in the stream) in Recognition and Free Recall. Here we study PI effects in RSVP.

Perhaps the simplest manipulation to study PI effects in RSVP is to manipulate the length of the study list and to compare performance in long lists relative to the short lists. It is expected that as more items are studied, the more difficult will be the retrieval of a particular item. The List Length Effect (LLE) has been consistently reported in recognition tasks (Cary & Reder, 2003; Ohrt & Gronlund, 1999; Strong, 1912) although other studies have reported a null LLE (Dennis & Humphreys, 2001; Kinnell & Dennis, 2011). Combining the length of the list with the positioning of the critical items in the long lists (critical items preceded by many items) enables one to explore PI effects. Therefore, in the present studies, we explored to what extent items presented on the fringe of awareness in RSVP can hinder the reportability of items encoded later in the list.

The task used in the present experiments is a classic old/new recognition question presented at the end of the stream. We recorded the accuracy and reaction times of the responses (e.g., Sternberg, 1975). The cognitive processes involved in this task have been investigated with Drift Diffusion Models (DDMs) (Ratcliff, 1978), which form part of a large family of models known as Sequential Sampling Models (Townsend & Ashby, 1983). These models assume that information supporting decision accumulates over time in a noisy process. Once the evidence supporting one choice or the other crosses a decision threshold, the response is executed. These models not only provide a close fit to the accuracy data but they also reproduce the Reaction Time distributions. DDMs assume a single accumulator integrating relative evidence over time until the evidence crosses one of two decision boundaries. The rate of accumulation is called the drift-rate (v). The distance between the two boundaries is called the threshold (a); it determines the amount of evidence that needs to be accumulated to produce a response. The threshold (a) is crucial to reproduce the common speed accuracy trade-off (SAT) (Schouten & Bekker, 1967), since the larger the threshold, the lower the probability of errors (lower probability that the crossing of a decision boundary is due to noise) and the larger the reaction times (more evidence needs to be accumulated to reach the boundary; Voss et al., 2004; Zhang & Rowe, 2014). Importantly, the threshold parameter *a* is determined by the participant and therefore it represents the subjective part of the model. Processes that are not encapsulated in the decision process (e.g., movement initiation and execution) are represented by the non-decision time parameter (*t*). Finally, the response bias is called *z;* it represents the starting point of the accumulation process relative to the boundaries.

In the current experiment, we explored potential PI effects on recognition performance for words presented at different rates (by manipulating the Stimulus Onset Asynchrony (SOA)) in RSVP. PI effects were induced by a manipulation of the length of the streams. Critical words were inserted in Short streams (12 words) or in the last 12 words of Long streams (36 words). At the end of each stream, a probe consisting of a New word or an Old word was presented. Participants were asked to indicate whether they saw the probe in the stream. In addition, accuracy and reaction times were used to fit the parameters of DDMs. A modulation of the parameters of the DDM, specifically the drift-rate *(v)* and the boundary *(a)*, can help to interpret the impact of PI in the recognition process. A modulation of the drift-rate *(v)* caused by the Length of the streams would shed light on the impact of PI on the memory processes outside the control of the subject. In contrast, an impact of the PI manipulation on the boundary parameter *(a)* will provide information about the subjective part of the decision process. This last part is particularly interesting since we are comparing conditions where stimuli are presented on the fringe of awareness (100ms of SOA), with supraliminal conditions (350ms of SOA).

In the following, we present two experiments to test our hypothesis that pro-active interference will be more evident in the long than short SOA conditions and that this will be particularly evident in DDM parameters. The second experiment represents a replication of the first.

2. Experiment 1

Method

Participants

20 (mean age= 19.55, sd= 1.16) undergraduate students (all female) of the University of Birmingham took part in this study in exchange for course credits. All were native English speakers and had normal or corrected-to-normal vision. The experiment conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Materials

192 English words were selected from the English Lexicon Project database (Balota et al., 2007) to serve as Old words and New words for the Recognition test. We selected 1236 words from (Warriner et al., 2013) as fillers to generate the RSVP streams. Two lists of Old-words were created and counterbalanced across stream-length conditions, while the same set of fillers was used in both lists. All words had 6 letters and none of them were proper nouns.

2.3 Procedure

We presented RSVP streams on a 24'' LCD screen (refresh rate: 60Hz, resolution: 1920 x 1080) using custom Psychopy3 (Peirce et al., 2019) scripts running under Python 3.6. Stimuli were Arial white characters of 1.3° of visual angle in height on a grey background (40% white). Participants were seated at 60cm from the screen. The experiment was conducted individually in a quiet room. Each experimental session consisted of 192 RSVP trials. 12 types of trials were generated by crossing the experimental manipulations (2 Length: Short, Long, x 3 SOA: 83, 100, 350ms x 2 Probes: Old, New).

Each trial started with the presentation of a fixation cross (+) at the centre of the screen for 1000ms. Immediately after the fixation cross, a sequence of words was presented followed by a mask (######). If the trial contained a short stream and the Old-word item was presented, this could be inserted in any position between 3rd and 9th. Crucially, if the trial contained an Old word in a Long-stream, this could be inserted in any position between 27th and 33rd (so the retention interval was similar in Short and Long streams). The trial ended with the question “Have you seen this word?” presented in the upper part of the screen along with a probe word presented at the centre. Participants were instructed to press the left-arrow key to indicate YES or the right-arrow key to indicate NO. If no response was registered in the 2s following the presentation of the probe the experiment continued with the next trial.

Table 1. Mean d’ values, mean RTs, and (standard deviations).

Length	d' score			Reaction Time		
	SOA			SOA		
	83	100	350	83	100	350
Short	0.53 (0.46)	0.79 (0.5)	1.53 (0.54)	0.833 (0.11)	0.826 (0.11)	0.808 (0.13)
Long	0.65 (0.4)	0.6 (0.52)	1.62 (0.62)	0.858 (0.13)	0.848 (0.12)	0.886 (0.14)

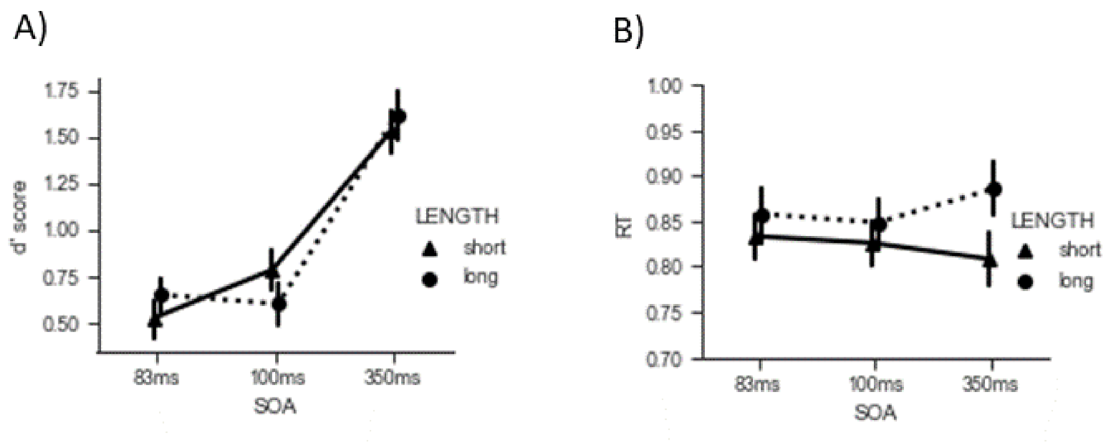

Figure 1. Main results of Experiment 1. d' scores (A) and Reaction Times (B) for Long (solid line) and Short streams (dotted line) and each SOA condition (horizontal axis).

Results

Reaction times and d'

We calculated the d' score from the Hit and False Alarm rates obtained in the responses to the Recognition question. A 3 (SOA: 83, 100, 350 ms) x 2 (Length: short, long) within subjects ANOVA revealed that the d' values (see Figure 1 and Table 1) were significantly larger for the longer SOA condition, $F(2,38)=53.72$, $MSE= 11.681$, $p<.001$, $\eta_p^2= 0.739$. Post hoc comparisons confirmed a significant difference between the SOA 350 and the SOA100, $t(19)=7.57$, $p<0.001$, $d=1.95$. Similarly, the d' score for the SOA 350 was significantly higher than the SOA83, $t(19)=10.49$, $p<0.001$, $d=2.53$. No significant difference was found between SOA 100 and SOA 83 ($p>0.1$). Neither the effect of Length nor the Interaction were significant (all $ps>0.1$).

A two-way repeated-measures ANOVA on the Reaction Times showed no significant effect of SOA ($p>0.1$). Longer Reaction Times were observed in the Long streams (mean= 0.86, sd= 0.13) relative to the Short streams (mean= 0.822, sd=012),

$F(1,19)=23.15$, $MSE= 0.052$, $p<0.001$, $\eta_p^2= 0.549$. Crucially, the interaction reached significance $F(2,38)=4.91$, $MSE= 0.010$, $p<0.05$, $\eta_p^2= 0.205$. Post hoc comparisons (Holm corrected) revealed that Long streams were not significantly different to Short streams in the SOA 83 condition ($p>0.1$). In contrast, Reaction Times in the Long stream were larger than those of the Short streams for the SOA 100 condition, $t(19)=2.48$, $p<0.05$, $d=0.19$ and for the SOA 350 condition, $t(19)=5.94$, $p<0.001$, $d=0.59$.

Hierarchical drift-diffusion model

The differences between conditions in the d' (higher d' for the longest SOA condition) and in RTs (an interaction between SOA and Length) indicated that there might be differences between conditions in the drift-diffusion parameters. To test this, we fitted a Bayesian Hierarchical Drift-Diffusion model to the accuracy and RT data. This method used Markov-chain Monte-Carlo (MCMC) to estimate the posterior distributions of DDM parameters. We used a model comparison approach to find the model that fitted the data best. We focused on the Drift-rate *v* and the Boundary separation parameter *a* to test the effect of the Length and SOA conditions. Non-decision time *t* and bias *z* were not allowed to vary between conditions, since no differences are expected in non-decision as the response display was the same for all conditions. Similarly, the bias *z* was not expected to vary in a non-blocked design, where trials of different conditions are randomly presented. Two models were compared: the Interaction Model that included as regressors the two variables SOA and Length and the interaction SOA*LENGTH for parameters *v* and *a*, and the Additive Model that included SOA and Length, for both parameters without the interaction. The result of the model comparison showed that the Interaction Model resulted in a better fit (lower Deviance Information Criterion), DIC = 3316, relative to the Additive Model, DIC = 3323. This result suggests that the interaction between SOA and Length contributes significantly to obtain a better fit to the data. To explore this interaction, we performed Bayesian hypothesis tests by comparing the probability mass of the posterior distribution for each of the 6 possible experimental conditions (3 SOAs x 2 stream lengths). Specifically, with this procedure we tested the amount of overlap of the posterior distributions of Long and Short streams for each SOA Condition. The results of the boundary parameter (see Figure 2) showed that the effect of Length only reached significance for the SOA 350ms $p<0.01$, while it was not significant in the 100 ms and the 83 ms conditions. The effect of length on the drift-rate parameter was not significant for any of the SOA conditions.

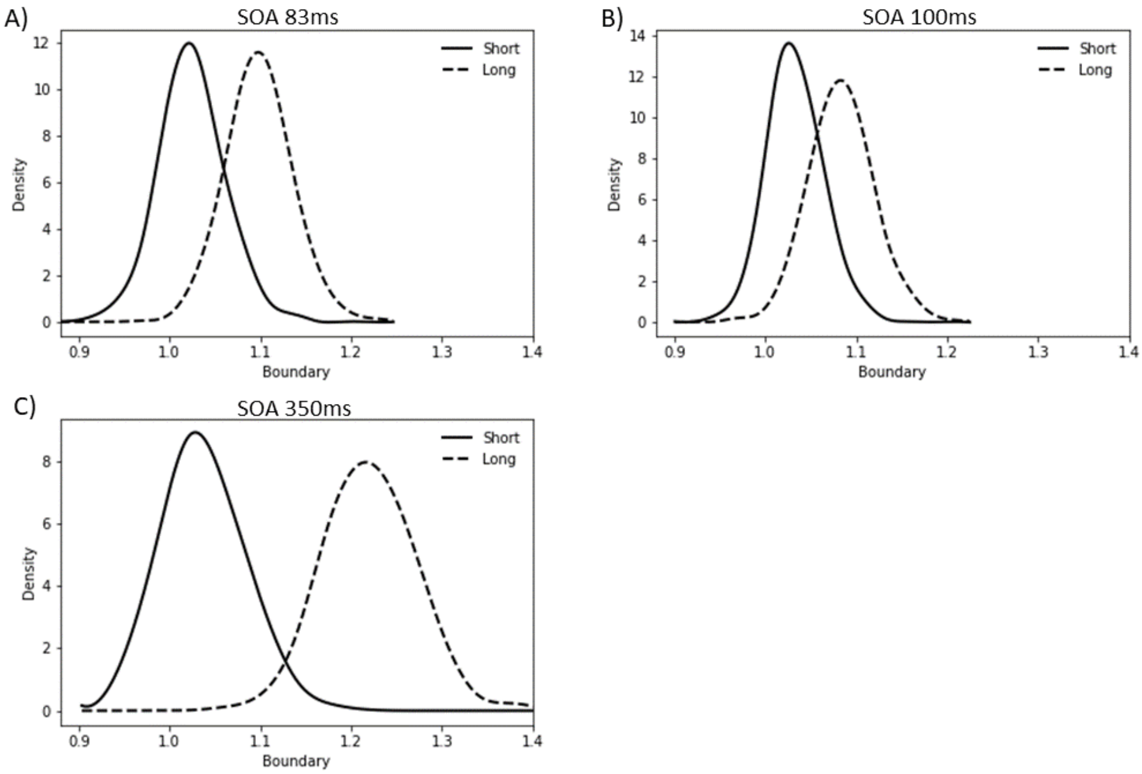

Figure 2. Posterior distribution of the Boundary parameter *a* for Short (solid line) and Long (dashed line) stream for the 83 ms (A), the 100 ms (B) and the 350 ms (C) SOA conditions.

3. Experiment 2

Method

Experiment 2 present a partial replication of Experiment 1. Experiment 2 was conducted online while Experiment 1 was run in the lab. In addition, Experiment 2 only tests the 83ms and 350ms of SOA.

Participants

47 (mean age= 19.43, sd= 1.23) undergraduate students (44 female) of the University of Birmingham took part in this study in exchange for course credits. All were native English speakers and had normal or corrected-to-normal vision. The experiment conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Power analysis

We performed a simulation based power analysis for an interaction using Superpower (Lakens & Caldwell, 2021) using the mean RTs and std of each condition (83ms and 350ms) of Experiment 1. To obtain a power of 90%, we needed 47 participants.

Materials

184 English words were selected from the English Lexicon Project database (Balota et al., 2007) to serve as Old words and New words for the Recognition test. We selected 1236 words from (Warriner et al., 2013) as fillers to generate the RSVP streams. Eight lists of Old-words were created and counterbalanced across conditions while the same set of fillers was used in both lists. All words had 6 letters and none of them were proper nouns.

Procedure

Participants were asked to use Google Chrome or Firefox as web browsers to run the experiment. We used Psychopy v2021.2.3 to code the experiment (Peirce et al., 2019) and it was hosted by Pavlovia.org. After the completion of the consent forms, hosted by Qualtrics (https://www.qualtrics.com), participants started the experiment by following the link provided at the end of the form. After a welcome screen, participants were instructed to work-out the size of the screen by adjusting a rectangle presented on the screen to the size of a credit card (Morys-Carter, 2021). Then, participants were presented with the instructions of the experiment. Each experimental session consisted of 184 RSVP trials. The structure and format of each trial were identical to experiment 1.

Results

Table 2. Mean d' values, mean RTs, and (standard deviations).

	d' score		Reaction Time	
	SOA		SOA	
	83	350	83	350
Short	0.79 (0.5)	1.51 (0.89)	0.911 (0.14)	0.851 (0.15)
Long	0.77 (0.55)	1.31 (0.68)	0.913 (0.16)	0.901 (0.16)

Reaction times and d'

We calculated the d' score from the Hit and False Alarm rates obtained in the responses to the Recognition question. A 2 (SOA: 83, 350 ms) x 2 (Length: short, long) within subjects ANOVA revealed that the d' values (see Figure 3 and Table 2) were significantly larger for the longer SOA condition. $F(1,46)=66.24$, $MSE= 18.725$, $p=0.000$, eta= 0.590 Bayes factor 10= 1.7×10^{13}. Neither the effect of Length nor the Interaction were significant (all ps>0.1).

A two-way repeated-measures ANOVA on the Reaction Times showed a significant effect of SOA F(1,46)=16.09, MSE= 0.062, p=0.000, eta= 0.259 BF10=8609. Longer Reaction Times were observed in the Long streams (mean= 0.86, sd= 0.13) relative to the Short streams (mean= 0.822, sd=012), F(1,46)=11.24, MSE= 0.033, p=0.002, eta= 0.196 BF10=165. Crucially, the interaction reached significance F(1,46)=13.60, MSE= 0.028, p=0.001, eta= 0.228 BF10=47. Post hoc comparisons (Holm corrected) revealed that Long streams were not significantly different to Short streams in the SOA 83 condition (p>0.1). In contrast, Reaction Times in the Long stream were larger than those of the Short streams for the SOA 350, t(46)=4.63, p<0.001, d=0.33 BF10=708.

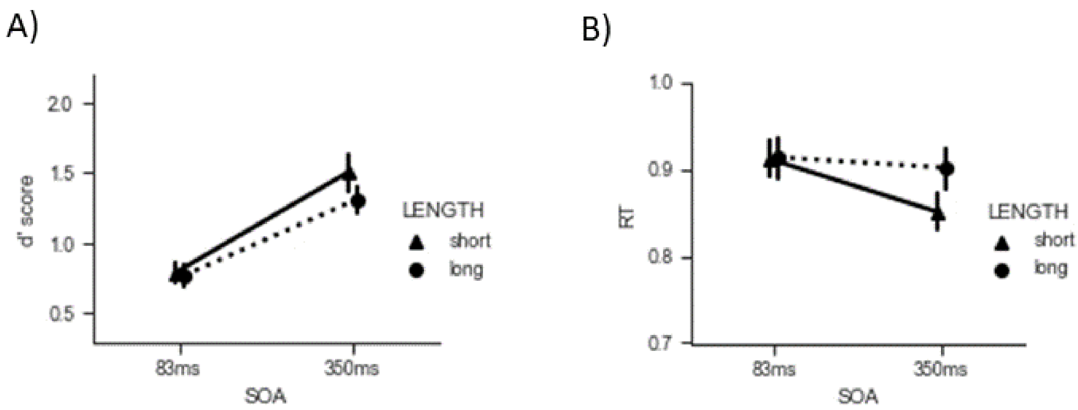

Figure 3. Main results of Experiment 2. d’ scores (A) and Reaction Times (B) for Long (solid lines) and Short (dotted lines) streams and each SOA condition (horizontal axes).

Hierarchical drift-diffusion model

Two models were compared: the Interaction Model that included as regressors the two variables SOA and Length and the interaction SOA*LENGTH for parameters *v* and *a*, and the Additive Model that included SOA and Length for both parameters without the interaction. The result of the model comparison showed that the Interaction Model resulted in a better fit (lower DIC), DIC = 9244, relative to the Additive Model, DIC = 9267. As in Experiment 1, the interaction between Length and SOA contributed significantly to obtain a better fit to the data. The posterior distributions of the *a* parameter (for each Length and SOA condition) showed significant differences between short and long streams in the slow streams, $p<0.001$; see Figure 4(B). Crucially, the same comparison was not significant in the fast streams; see Figure 4(A). The difference between Short and Long streams for the drift-rate approached significance for the Slow streams $p=0.095$; see Figure 4(D), and was not significant in the Fast condition $p>0.1$; see Figure 4(C).

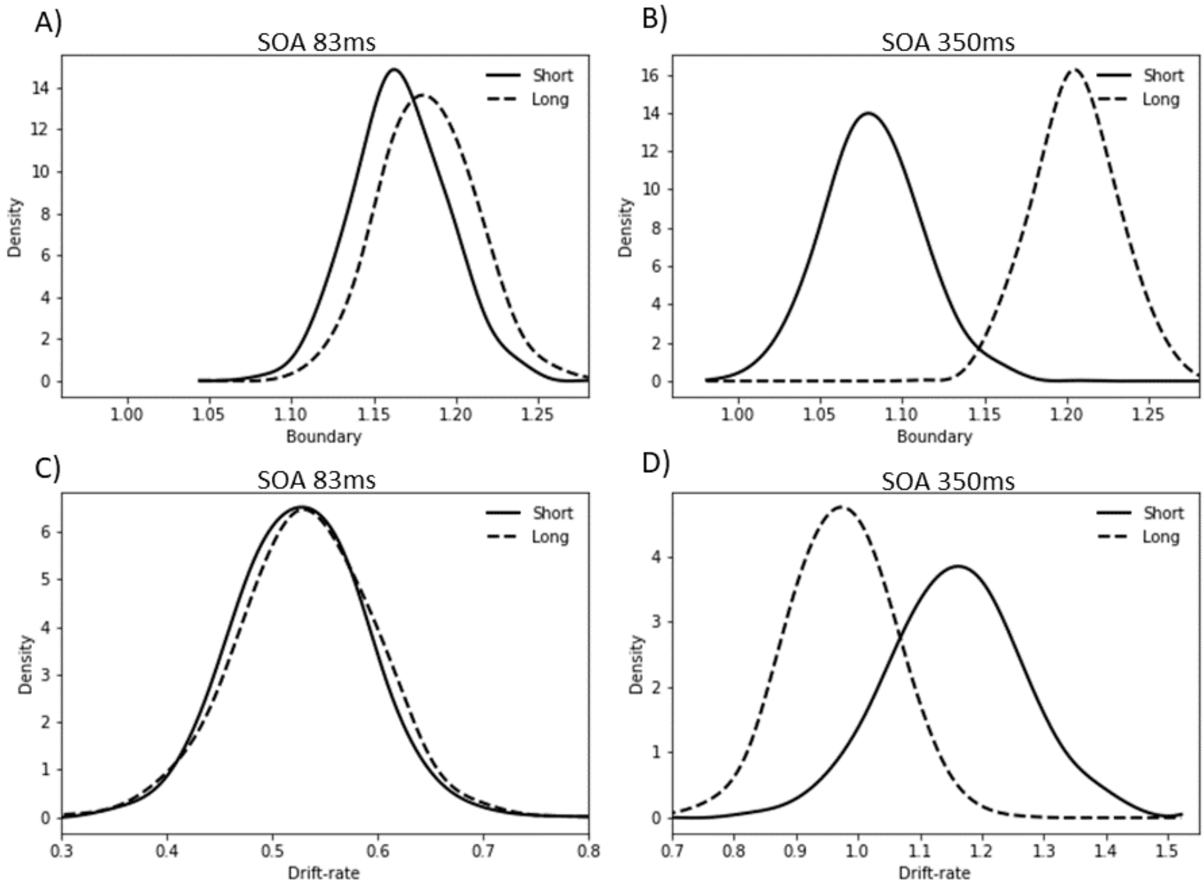

Figure 4. Posterior distribution of the DDM in Experiment 2. A) parameter *a* for the 83 ms SOA, B) parameter *a* for the 350 ms SOA, C) parameter v for the 83 ms SOA, D) parameter v for the 350 ms SOA. Each panel contrasts the Short (solid line) with Long (dashed line) streams.

4. Discussion

In two experiments, we compared the LLE for words embedded in RSVP streams at different SOA conditions (83, 100 and 350ms in Exp. 1 and 83 and 350ms in Exp. 2.). We used a proactive interference design, with the critical items of the long list presented at the end of the list, and we obtained d' values and reaction times from a single probe recognition task. We found a null LLE on the d' values in every SOA condition. In contrast, a LLE was shown in the reaction times of the slow streams, with longer reaction times in the long lists relative to the short list. In addition, we investigated this effect on the drift-rate and the boundary parameters of a Hierarchical Bayesian Drift Diffusion Model fitted to the reaction times and accuracy data. The list length manipulation did not result in a modulation of the drift-rate in any SOA condition. In contrast, we observed larger boundary values in the long lists of the slow streams. To our knowledge this is the first study to investigate the LLE on reaction times and the drift-diffusion parameters of a recognition task for words embedded in RSVP streams.

A LLE in Recognition has been consistently reported in the literature (Cary & Reder, 2003; Ohrt & Gronlund, 1999; Strong, 1912), with some notable exceptions (Dennis & Humphreys, 2001; Kinnell & Dennis, 2011). The evidence from this study is consistent with a null LLE on accuracy/discrimination data. The list length manipulation did not modulate the d' values regardless of the speed of the streams. Kinnell and Dennis (2011) conducted a series of experiments investigating the possibility that certain confounds could be the cause of the discrepancies in the literature. One of these confounds is the use of a proactive design, namely the presentation of the critical items (those that could be probed later) by the end of the long list. It is argued that the differences in performance between long and short lists are caused by differences in attention. It is less likely that participants sustain attention during the entire block in long than in short lists. The LLE would be caused by this attentional depletion of the long list rather than the noise introduced by the large number of items encoded. Interestingly, the present study employed a proactive approach but the expected LLE -due to attention or interference- on the d' was not found. We believe that this is not really surprising once one considers that the "attention confound" is negligible in the present studies. The long lists in this study lasted 12.6 seconds, while the long lists in a classical recognition experiment, where words are presented for 2 or 3 seconds, could last for hundreds of seconds. The duration of a long list in the present study is equivalent to the first 4 or 5 words in a typical recognition experiment. The modulation of attention within trials is not a sizeable confound in the present experiments, and, therefore, the null LLE is in line with the Kinnell and Dennis' findings.

The null LLE reported here is consistent with the predictions of context-noise models of recognition memory. For example, in the BCDMEM model (Dennis & Humphreys, 2001) Recognition performance depends on the degree of match between the context retrieved at test and the study context. Any interference observed should come for the previous context where the words have been encountered and the interference from other words in the same list is negligible. In contrast, in noise-item models such as MINERVA II (Hintzman, 1986) performance is mainly drive by the match between the probe and the memory traces for items in the list. In this family of models, all items play a role in the computation of the degree of overlap and generally the more items presented, the larger the noise and, consequently, a poor performance is expected after long lists. However, although the present findings are consistent with the prediction of context-noise models, we believe that these experiment do not constitute a direct test for models that were not developed to account for the data obtained with RSVP.

In contrast to the d' results, we found an LLE on the RTs of the recognition test. However, this does not necessarily imply that there is a difference between long and short list on retrieval speed. As pointed out by Öztekin and McElree (2007), RTs are sensitive to multiple factors: retrieval speed, the quality of the match between the probe and the memory representation, or simply a difference in the criterion set by the participants. To infer a difference in retrieval speed caused by proactive interference, Öztekin and McElree (2007) used a Speed Accuracy Tradeoff variant of a probe recognition task. They found that proactive interference from items in previous trials affected the growth of accuracy (participants were cued to respond at different intervals) over retrieval time. They concluded that proactive interference had an impact on the retrieval of words in a recognition task. In contrast, the LLE observed on the RTs in the present study does not point to a conclusion based on retrieval speed. The parameters of the drift diffusion model showed that there were no significant differences in the drift-rate, which is the parameter that more closely could be interpreted as reflecting the speed of information accrual (Ratcliff, 1978). Therefore, there is little evidence in the present study to support the position that list length manipulation is having an impact on the retrieval speed or the match between probe and memory representations, since there are similar d's between long and short lists.

An alternative interpretation is that the longer RTs in the long lists of the slow streams are caused by participants applying a more conservative criterion in those trials. The values of the boundary parameter of the fitted DDM seem to be consistent with this interpretation. Interestingly, this LLE on the boundary parameter was not observed in the fast RSVP streams. That is, in the fast streams, participants were as conservative in the long as in the short streams. The extra 24 words presented in the long lists of the fast streams did not have an impact on the conservatism of the participants. This is most noticeable in the absence of a LLE for the fast streams on the boundary parameter of the model. Thus, there is no LLE on the participants objective performance (d', reaction times and drift-rate parameter) nor on the subjective decision (boundary parameter of model) for items presented in RSVP.

As stated in the introduction, we are interested in testing the possibility of episodic memory encoding without conscious perception of the stimuli. The results of the fast streams seem to be consistent with the absence of episodic encoding of most of the items in RSVP. The additional presentation of 24 items did not have an impact on the accuracy, the retrieval speed, nor on the criterion placement. Particularly, the absence of a modulation of the criterion seems consistent with many of the words not being consciously perceived. The presentation of these additional words failed to modulate the subjective aspect of the recognition task.

In earlier works previous chapters, we have presented an account of recognition data in RSVP suggesting that most of the items presented in RSVP fail to be encoded in long-lasting episodic representations. We argue that this is in part because many items in RSVP remain below the threshold for conscious perception. The present findings are consistent with this interpretation. If most of the items in RSVP are not episodically encoded, then the degree of interference expected in a long list of words might be negligible. In agreement with this hypothesis, the present findings suggest that in the fast RSVP streams no form of interference is observed: neither in the sensitivity nor in the criterion set by the participants. This is a good fit for The Tokenized Percept Hypothesis, Chapters 1 and 2, that suggests that the long lasting episodic encoding is a feature of conscious processing. It is suggested that Type information (perceptual and categorical attributes) is activated regardless of the awareness of the stimuli. The activation of Type information would support the cases of perception without awareness (e.g. masked priming, Ferrand, 1996) and would be responsible for the high accuracy in detection and identification tasks in RSVP. In contrast, Token information (the attribution of Type information to a particular temporal context) requires a second stage of processing where representations targeted by attentional processes are bound to episodic features, such as, information about the order or the position relative to other stimuli in a sequence. Many of the items presented in RSVP would not reach this second stage of processing and therefore they would not acquire explicit episodic properties. These missed items would not interfere with the retrieval of subsequent items and, not being consciously perceived, they would fail to impact the subjective element of the recognition task. This Tokenized Percept Hypothesis takes inspiration from the Simultaneous Type/ Serial Token model (Bowman & Wyble, 2007; Bowman et al, 2008); it also has similarities to a theory of perceptual experience introduced by Kanwisher (Kanwisher, 2001).

Chapter 6 extends the investigation into the impact of early presented items on the reportability of stimuli in Rapid Serial Visual Presentation (RSVP). Specifically, the focus is placed on examining the repetition effect, akin to the study conducted in Chapter 2, which explores the potential facilitation in processing items with identical identities. However, in contrast to Chapter 2, the repetitions are integrated within the same stream, while manipulating the number of intervening distractors presented between different instances of the identical item.

Chapter 6: The spacing effect on the buildup of memories in RSVP

1. Introduction

Our environments are full of regularities that shape our behaviour and cognition. Often these regularities are experienced as mere repetitions, and it is through repetition that they are consolidated as memories and made resistant to forgetting. But we do not notice *all* the repetitions, simply because we cannot register consciously the vast majority of the stimuli at any given time. These missed repetitions can still be registered without reaching awareness, but it is not clear to what extent they contribute to our learning and the consolidation of memories. Furthermore, these repetitions might be too spaced (or too close) to induce a discernible effect. That is, there is a spacing effect that modulates memory consolidation (Cepeda et al., 2006; Verkoeijen et al., 2005). In this paper, we will explore the contribution of missed repetitions on the consolidation of memories and how this contribution is modulated by the distance between presentation of repeated stimuli.

Repetition detection

Our perceptual and memory systems possess an immense capacity for detecting repetitions even when they are presented very rapidly. Thunell and Thorpe, (2019), embedded repeated images in Rapid Serial Visual Presentations (RSVP) among distractors in a repetition detection task. They found above chance performance after two presentations and large improvements in performance as the number of presentations increased. They concluded that there is a mechanism in the early stages of processing for detecting repetitions. This mechanism would play a crucial role in learning (Agus et al., 2010).

The nature of the proposed mechanism -prior to working memory- and the fact that this impressive performance was obtained with RSVP stimuli might suggest that the conscious/unconscious status of stimuli is irrelevant for being detected as repeated. In other words, all presentations of a stimulus contribute similarly to the repetition detection performance. This is consistent with the observed high performance after many repetitions as it suggests that the strength of the memory trace increases across repetitions. However, there is an alternative interpretation that questions the nature of the evidence accumulation process (Avilés et al., 2020; Bowman & Avilés, in press). Higher than chance performance in the repetition detection task might require conscious perception of at least some of the repetitions. That is, to detect a repetition, the conscious processing of two presentations is sufficient. In paradigms like RSVP, more presentations of the same stimulus simply gives more chances for two presentations to be consciously perceived. In this account, there is no evidence accumulation across presentations that are not consciously perceived.

The contribution of missed repetitions to the overall repetition detection performance is the key factor that differentiates these accounts. On the no subliminal evidence accumulation account, the memory traces of missed repetitions are too weak to have a noticeable impact on subsequent presentations. Thus, the probability of detecting a repetition *for the first time* across presentations should be constant. Aviles et al., (2020), showed that when instances of the same item are presented in different RSVP trials, there is an invariance to the number of repetitions in the probability of first detection as a repetition. They concluded that there is no evidence accumulation across subliminal repetitions. Most of the items would fail to be consciously perceived and therefore many repetitions would be missed. Repetition detection performance would be based on the few items that by random factors (e.g., poor masking) are consciously perceived.

In Aviles et al., (2020) there were long gaps between repetitions. It could be argued that for evidence to accumulate, perhaps a shorter gap is required so that the memory trace does not vanish (due to decay or interference from distractors). There is some initial evidence in favour of the modulating factor of spacing in the build-up of memories in RSVP. Bowman and Aviles (2022), analysed Thunell and Thorpe (2019)'s data where presentations were separated just by one or two distractors. They showed that in the two-distractor condition there, was no evidence accumulation. That is the probability of detecting a repetition for the first time was invariant to the number of repetitions. However, the results were inconclusive in the one-distractor condition (Bayes Factor 1.83). These results open the possibility that subliminal evidence accumulation could be observed when repetitions are sufficiently close. That is, even if not consciously perceived residual activation following the presentation of a target would last long enough to impact subsequent presentations.

Recognition memory

Repetition plays a fundamental role in the consolidation of memories. It has been long known that repeated events are less likely to be forgotten (Ebbinghaus, 1964). As in the case of repetition detection, performance in a recognition task improved with the number of presentations (Endress & Potter, 2014). However, the effect of spacing is more debatable. While in the memory literature it has been consistently reported that spaced repetitions improve performance relative to massed presentations, in RSVP the manipulation of the gap between presentations has yielded mixed results. In Martini and Maljkovic, (2009), the spacing effect was not found, while Thunell and Thorpe (2019) found a spacing effect in an identification task presented after each trial. Martini and Maljkovic, (2009), argued that memory performance followed the total time hypothesis that states that the main factor that modulates recognition performance is the amount of time the items have been presented.

Another crucial aspect of the relationship between repetition and memory consolidation concerns the role played by the explicit detection of repeated items. As stated above, Endress and Potter (2014) showed that performance in a recognition test for images presented in RSVP was better the more times the items have been presented. This, again, suggests a kind of evidence accumulation even for rapidly presented items at the threshold of awareness. Similarly to the case of the repetition detection, there are two alternative interpretations for this effect. If every item, regardless of being detected as a repetition, contributes to the memory performance then even items not detected as repeated would present a robust repetition effect. There is some initial evidence that contradicts this hypothesis. The repetition effect for items not detected as repeated in Thunell and Thorpe (2019) is negligible. Here we provide more evidence in the same direction.

One way to gain more insight about the relationship between consciousness and the processing of repetitions is to look at metacognitive processing. Not only is it widely accepted that conscious content is generally available for introspection, but there is also evidence that conscious processing and metacognition could share some brain mechanisms (Lau & Passingham, 2007). By asking participants to rate their confidence about the accuracy of their answers to the recognition question, one can obtain information about the efficiency of the metacognitive processes in tracking the recognition decisions. One can calculate the meta-d’ (Maniscalco & Lau, 2012) to infer the information available for metacognition. By tracking the metacognitive efficiency, the gap between meta-d’ and d’, one can extract conclusions about how much information is not accessible to metacognition. It is to be expected that the d’ of the recognition performance will improve with the number of repetitions. If there is evidence accumulation that is below threshold -not available for metacognition- then the improvement of the meta-d’ would be less pronounced and the gap between d’ and meta-d’ would increase with the number of repetitions, signalling a loss of metacognitive efficiency.

In the present experiment, we present repeated items embedded among distractors in RSVP. The critical item could be presented from 1 to 8 times. Crucially, we will use words as in Aviles et al., (2020). The use of words instead of images will allow us to explore the generalizability of Thunell and Thorpe (2019) findings to a different category of stimuli. Participants will be asked to search for repetitions. At the end of every trial, they will be presented with a 2-alternative forced-choice recognition question, a confidence question in the accuracy of the answer to the recognition question and a repetition detection question.

2. Method

Participants

58 (mean age=19.04, sd =0.89) undergraduate students (52 female) of the University of Birmingham took part in this study in exchange for course credits. All were native English speakers and had normal or corrected-to-normal vision. The experiment conformed to British Psychological Society criteria for the ethical conduct of research and ethical procedures of the School of Psychology at the University of Birmingham.

Materials

480 English words were selected from the English Lexicon Project database (Balota et al., 2007) to serve as probes for the 2-alternative forced-choice recognition test. Half of the words had 5 letters and the other half had 6 letters. Half of the probes for the recognition test served as Critical items in the Repetition conditions. We selected 2490 words (44% had 5 letters and 56% had 6 letters) from (Warriner et al., 2013) as fillers to generate the RSVP streams.

Procedure

Participants were asked to use Google Chrome or Firefox as web browsers to run the experiment. We used Psychopy v2021.2.3 to code the experiment (Peirce et al., 2019) and it was hosted by Pavlovia.org. After the completion of the consent forms, hosted by Qualtrics (https://www.qualtrics.com), participants started the experiment by following the link provided at the end of the form. After a welcome screen, participants were instructed to work-out the size of the screen by adjusting a rectangle presented on the screen to the size of a credit card (Morys-Carter, 2021). Then, participants were presented with the instructions of the experiment. They were informed that their task was to search for a repeated word in each stream/trial. Each experimental session consisted of 240 RSVP trials. 72 trials contained a single presentation of the Critical item and the rest 168 trials consisted of 21 types of trials generated by crossing the experimental manipulations (1 to 3 Distractors x 2 to 8 Number of Presentations).

Each trial started with the presentation of a fixation cross (+) at the centre of the screen for 1000ms. Immediately after the fixation cross, a sequence of words was presented followed by a mask (######). Each RSVP stream contained 40 words. The SOA was set at 117ms. The last presentation of the Critical item was inserted in a random position between the 31st and 38th item. For trials containing repetitions, each repetition was separated by a number of distractors depending on the condition. The trial ended with the 2-alternative forced-choice recognition question "Which word has been presented?" in the upper part of the screen and the presentation of the 2 alternatives, the Critical item and a new word (not presented in the stream), in the centre-left and centre-right positions. After this, they were presented with the confidence question "How confident are you that your previous answer is correct?" in the upper part of the screen and a 6-point scale (from "not sure" to "very sure") in the centre of the screen. Participants were instructed to type the corresponding number. Finally, the trial ended with the presentation of the repetition question "Have you seen a repeated word?". They were instructed to respond "yes" or "no" by pressing the left or right arrow of the keyboard.

Results

The Repetition Detection, the Evidence Accumulation and the Recognition results are reported below. The Metacognition data (meta-d') have not been analysed and they are not included.

Repetition Detection

We used a linear mixed model (with Subject as Random Factor) to test the effect of Number of Presentations on the accuracy data for each Distractor condition. Accuracy improved as the Number of Presentations increased as indicated by the positive slopes of the models, see Figure 1. Distractor 1: coefficient=0.077, $z=16.62$, $p<0.001$. Distractor 2: coefficient=0.060, $z=12.17$, $p<0.001$. Distractor 3: coefficient=0.047, $z=8.77$, $p<0.001$. To test whether performance was above chance, we obtained the false alarm rate from the trials without repeated words, i.e. reports of a repetition, when there was no repeating item. Accuracy was above chance for all conditions, except when the Critical Word was presented twice separated by two distractors (the 95% CIs were above the false alarm rate, 0.284, for Distractor 1 and 3 conditions: see Figure 1). From the 4th presentation of the Critical Word onward, performance was clearly better when the presentations of the Critical Word were separated by one distractor relative to 2 and 3 Distractors (see Figure A.3.1 in the Appendix for the 95% CIs for the difference between conditions).

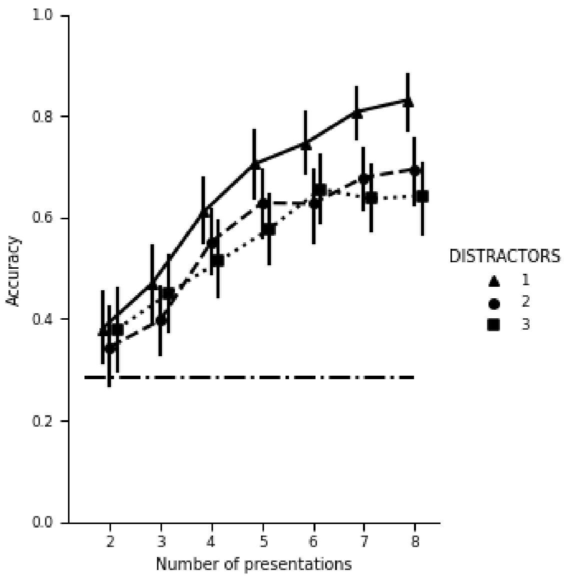

Figure 1. Main result from the Experiment. Accuracy at detecting repetition across number of presentations for each of the Distractor conditions. The dashed horizontal line represents the False Alarm rate, i.e. the probability of reporting a repetition on the first presentation, before there has been a repetition. Error bars indicate 95% confidence intervals.

Subliminal Evidence Accumulation

To test the presence or absence of subliminal evidence accumulation, we compared the accuracy data with the expected performance under the assumption of no subliminal evidence accumulation. This is a different approach to assess the same question as that tested in Aviles et al., (2020). Specifically, in Aviles et al., (2020) we quantified the probability of first seeing as a repetition and showed it did not change across numbers of presentations, suggesting no subliminal evidence accumulation. However, here we determine how the probability of seeing as a repetition (whether for the first time or not) would increase across presentations if the probability of *first* seeing was constant. That is, the absence of subliminal evidence accumulation was used as the Null Hypothesis. In order to generate the expected data under the Null, we reasoned that if there is no evidence accumulation then all missed repetitions prior to repetition *j* should have no effect on the accuracy increments. Which is to say that every increment of performance between presentation *j* and presentation *j+1* should be proportional to the difference between the baseline (false alarm rate) and presentation *2* (since this is actually the case where missed repetitions have no effect because there are just two presentations). Assuming this, we can generate the Null Hypothesis data just with the presentation *2* accuracy data and the difference between baseline and presentation *2*.

$$Accuracy_{(i+1)} = Accuracy_{(i)} + \frac{\left(Accuracy_{(2)} - Baseline\right)}{(1 - Baseline)} * \left(1 - Accuracy_{(i)}\right)$$

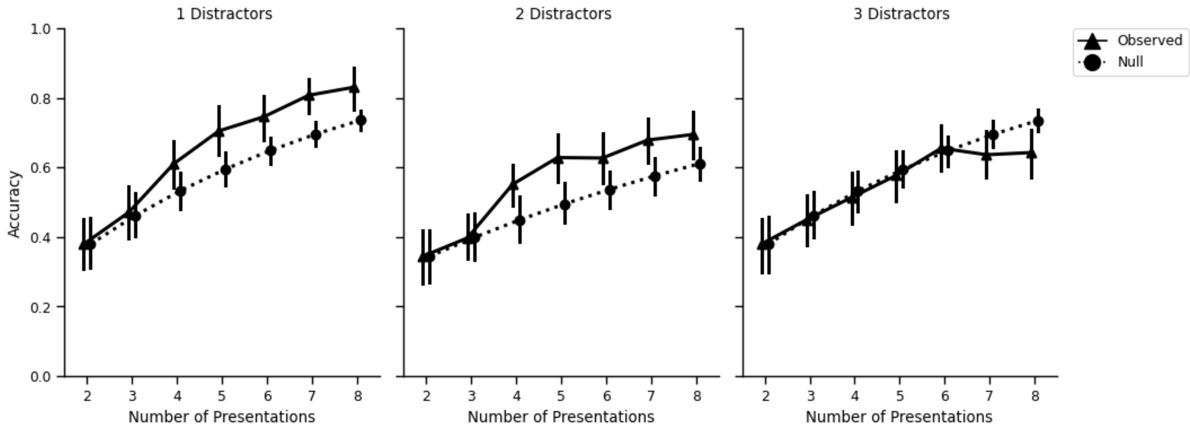

Figure 2. Evidence accumulation test. Accuracy (solid line) and Null Hypothesis data with no evidence accumulation (dashed line) for the one-distractor condition (left), two-distractors (centre) and three-distractors (right). Error bars indicate standard errors of the mean.

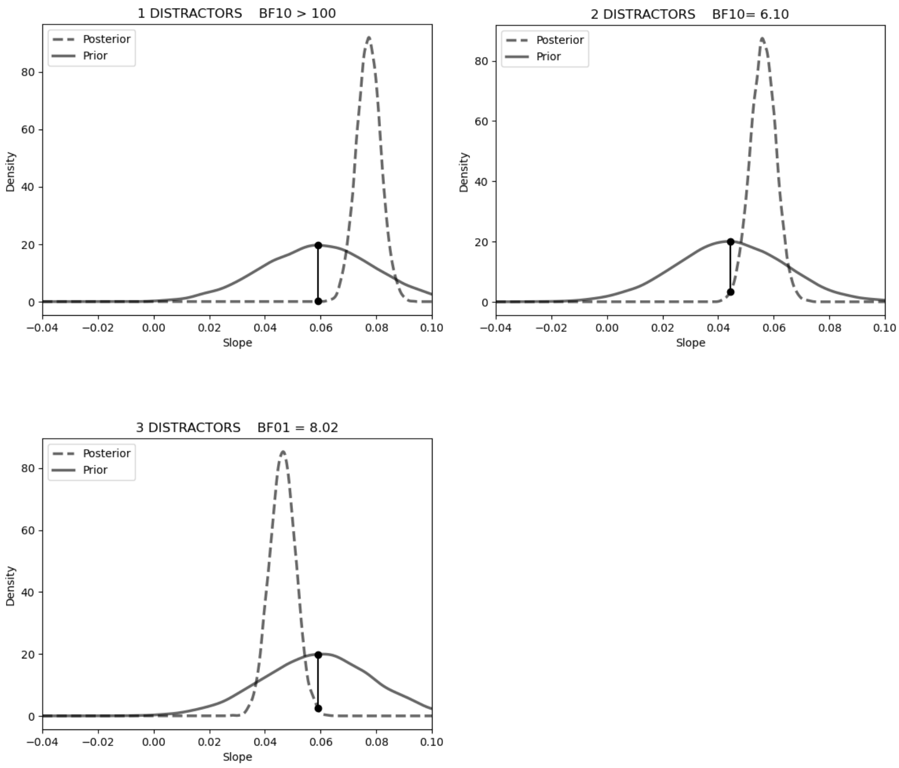

Figure 3. Posterior (dashed line), prior (solid) and Bayes Factor (density ratio)(vertical line at point of interest), for the one-distractor condition (top-left), two-distractors (top-right), and three-distractors (bottom-left).

To test whether the results (see figure 2) support the Alternative Hypothesis of subliminal evidence accumulation or the Null Hypothesis, we used a Bayesian regression mixed model (with the Bambi library in Python, Capretto et al., 2022). We obtained the posterior distribution of the slope of the model (see figure 3) and Bayes Factors were calculated using the Savage-Dickey method (Wetzels et al., 2009). This method consists of calculating the ratio between the density of the posterior and the prior at the point of interest. The point of interest in our case is the coefficient for the slope of the Null Hypothesis data. Using this method for the 1 Distractor condition we obtained that the data provided support for the Alternative with a BF10 > 100. Similarly, we observed that there is evidence for the Alternative (BF10 = 6.1) in the case of the 2 Distractor condition. In contrast the results for the 3 Distractor showed that the slope of the observed data was lower than the slope of the Null Hypothesis data (BF10 = 8.02).

Recognition

Data from the 2-alternative forced-choice question and the confidence question were fitted with a Hierarchical Bayesian Regression Model (with the ghsdtr library in R, Paulewicz and Blaut, 2020) to obtain the posterior distributions of the d' for each distractor condition. We used Bayesian Hypothesis testing (Kruschke, 2013) to test the significance of the effects. To generate *p*-values, we calculated the proportion of samples where there is overlap between two posterior distributions. Unsurprisingly, Critical Words that have been detected as repeated were recognized better than critical words not seen as repeated (all ps<0.001); see figure 4. There were no significant differences between distractors when the words have been detected as repeated (all ps>0.1). In contrast, recognition performance was poorer for the 3 distractor condition relative to the

1 Distractor ($p<0.05$ at presentations: 5, 6, 7) and 2 distractor condition ($p<0.05$ at presentations: 3, 4, 5, 6, 7) where words were not seen as repeated. For words seen as repeated, recognition performance improved with the number of presentations. For every Distractor condition, we compared the posterior at presentation 2 separately with each of the remaining Numbers of Presentations (i.e. 2 against 3 and then 2 against 4 and so on) conditions. For words detected as repeated, we observed that the overlap in Density Intervals decreases as the Number of Presentations increases and a significant difference was reached by Presentation 5 and this was the case for all spacings (1-Distractor, 2-Distractor and 3-Distractor). In contrast, for words not detected as repeated, the improvement of performance with the number of presentations only reached significance for the 1 Distractor condition (presentations 6, 7 and 8 were significantly different to presentation 2). In the 2 and 3 Distractor conditions, performance did not improve significantly.

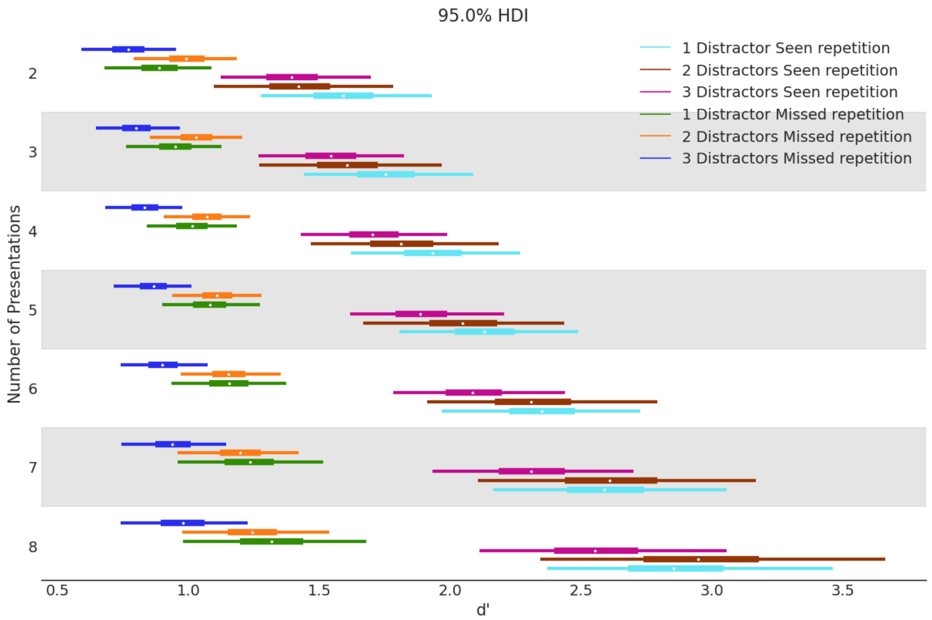

Figure 4. Main results recognition task. 95% High Density Interval (HDI) of the posterior distribution of d' values (horizontal axis) for each Distractor condition for Seen and Missed repetitions across presentations (vertical axis).

3. Discussion

We investigated the repetition and spacing effects in a repetition detection and a recognition task. As expected, performance in the repetition detection task improved with the number of repetitions. We found a clear spacing effect, with participants being more accurate in the one-distractor condition relative to the two and three-distractor conditions; see figure 2. These results confirmed the findings reported in Thunell and Thorpe (2019) in a similar study with pictures instead of words as stimuli. Crucially, the improvement of performance on the Repetition task with the number of presentations was modulated by the number of distractors between presentations. Larger improvements were observed in the one-distractor condition relative to the two and three-distractors conditions; see figure 2.

Do all repetitions contribute equally to the repetition detection decision? We investigated whether this process of evidence accumulation takes place even while repetitions are not detected, i.e. subliminally. To test this, we compared the repetition detection performance (solid lines in figure 2) with the hypothetical performance in the absence of evidence accumulation for missed repetitions (dotted lines in figure 2). We found that there is a clear pattern of (subliminal) evidence accumulation for the one-distractor and two-distractor conditions. In contrast, there was no evidence accumulation for the probability of first detection of a repetition in the three-distractor condition. When a similar analysis was conducted on Thunell and Thorpe data (Bowman and Avilés, 2022), we found an absence of evidence accumulation in the two-distractor, but the data were inconclusive (Bayes Factor 1.83) in the one-distractor condition. The absence of a clear (subliminal) evidence accumulation in the one-distractor condition in the Thunell and Thorpe data might be explained by the fact that there is a ceiling effect in their data so not all data points could be used to calculate the probability of first detection rendering a less sensitive statistical test. However, it is more relevant that in both studies, there is an inverse relationship between the gap between presentations and (subliminal) evidence accumulation. As the gap increases from one to three distractors, the (subliminal) evidence accumulation vanishes. This is consistent with Aviles et al., (2020), that reported an absence of (subliminal) evidence accumulation when repetitions were inserted in different trials, and therefore separated by a large number of distractors.

It has been reported that recognition performance for unconsciously perceived items is extremely low (Chong et al., 2014). To the extent that the repetition detection task recruits similar processes as a recognition task then we can conclude that the high performance in the repetition detection task is based on consciously perceived stimuli. However, it is more difficult to argue that missed repetitions are also conscious. First, it would be surprising that participants fail to detect as repeated items stimuli that could be reported. Second, previous studies consistently agree that repetitions facilitate the processing of the second presentation of an item but only when the first occurrence could not be reported (Dux et al., 2006; Kanwisher, 1987; Shapiro et al., 1997). Dux et al., (2006), suggested that a Type/token account of the effect could explain the facilitation observed when two items in an RSVP stream share the same identity. In the Type/Token account, information about the featural and categorical information of stimuli are processed separately from the spatiotemporal information. Failures to bind a second token to an activated type result in Repetition Blindness (Kanwisher, 1987), viz, the non reportability of a second occurrence of an item. In this account, repetition facilitation is explained by the activation of type information from the first occurrence of a repetition and a failure to bind a token to a first occurrence of an item. The residual activation of type information would facilitate the encoding of subsequent occurrences.

Recognition

Recognition performance was much better in trials in which the critical item has been detected as repeated; see figure 4. This suggests that both tasks, repetition detection and recognition, recruit similar processes. More interestingly, the repetition effect is more pronounced when the words have been detected as repeated. Endress & Potter, (2014), reported a similar effect for recognition tests after a block of trials. The present findings suggest that the repetition effect on recognition depends on the explicit detection of the items as repeated. This is consistent with what we argued in Aviles et al, (2020), that, without the detection of a repetition, the improvements with the number of presentations in a recognition test would be much lower. The negligible small effect of repetition for missed repetitions confirms our interpretation. Alternatively, this increase could arise because Missed repetitions could still contain a single percept of the critical item and the probability of this single percept may increase with the number of presentations.

In summary, taken together, the repetition detection and the recognition results, suggest that there is a clear repetition effect, the more presentations, the better the performance. However, undetected and detected repetitions seem to have very different profiles that could indicate that the conscious perception of stimuli plays a crucial role in episodic encoding. Undetected repetitions facilitate the processing of subsequent presentations only when the gap between presentations is less than 300ms. As stated above, in the type/token framework, an undetected repetition that result in the facilitation of a second occurrence might be explained by a failure to bind a spatiotemporal token to the activated featural (type information) representation. The residual activation of the identity of the item could facilitate the processing of the second occurrence. The boost of Type activation could facilitate the conscious encoding of a subsequent presentation and, consequently, this could increase the likelihood of being detected as repeated; i.e. we would be observing subliminal evidence accumulation. Interestingly, even in the three-distractor condition (where no subliminal evidence accumulation is observed) performance is clearly above chance. This means that stimuli that do not benefit from the boost from earlier presentations can still be consciously perceived and be detected as repeated, even when they are separated by many distractors, as was observed in Aviles et al., (2020).

In Aviles et al. (2020), we introduced the Tokenized Percept Hypothesis, an account about the role of consciousness in the episodic encoding of stimuli. In this view, conscious perception happens when a Type, representing featural and categorical information, is bound to a token that encapsulates its episodic properties. In our account, the repetition task demands fully-fledged conscious percepts, since it requires access to the episodic information about items. We think that many of the items in RSVP fail to reach the threshold of awareness and will not benefit from the episodic durability that tokens provide. However, most of the items in RSVP streams would activate the corresponding type information (e.g., lexical representations in the case of words) that could facilitate the processing of posterior stimuli with similar characteristics. We think that this type of repetition priming is the source of the evidence accumulation reported here.

Chapter 7: General Discussion

The main objective of this thesis was to explore to what extent items presented in RSVP are encoded as episodic representations. We have presented evidence consistent with two conclusions: a) many RSVP items fail to be reported, b) the repetition, recency and proactive interference effects observed in RSVP suggest that the majority of items failed to be encoded as episodic representations.

1. Summary of findings

In the Introduction, we reviewed part of the evidence and hypotheses exploring the possible Functional Correlates of Consciousness (FCC). We demonstrated that cognitive processes that can only be carried out with consciousness have proven elusive. This has promoted the view that the FCCs simply do not exist (Hassin, 2013, Rosenthal, 2008), or that consciousness merely facilitates some cognitive processes (Shea & Frith, 2016).

The last segment of our review revisited an effect commonly observed in RSVP, but that has been obscured by the exuberant productivity of the AB and the RB paradigms. The gap, the "missing memory imprint", between recognition (reportability) and detection performance in RSVP (Potter, 1976). We think that this effect is suggestive of mechanisms that are highly efficient in the absence of conscious reports. Chapter 4 explored in depth the non-reportability of items in RSVP. We attempted to remove two possible confounds for any experiment exploring this issue. First, we reduced the retention interval between the critical item and the probe by testing the last item (before the mask) of the sequence. Second, we tested the reportability of task-irrelevant items in

RSVP by creating conditions for incidental encoding: the main task for participants was a detection task, while the memory for distractors was tested with a surprise recognition question. Our Bayesian analysis confirmed that the results provided evidence in favour of the hypothesis that performance was at chance level (d' = 0). This is consistent with the claim that RSVP task-irrelevant items leave negligible memory traces. Furthermore, this observed chance-level performance suggests that the "missing memory imprint" is even higher for distractors in RSVP. We confirmed this by comparing the recognition performance in the "surprise experiment" with that observed in an experiment where all items were task-relevant. Unsurprisingly, the results showed a significant difference between the two experiments. In conclusion, the evidence reported in Chapter 4 is the most "direct" evidence in support of the non-reportability of task-irrelevant items in RSVP, at least when conditions (with some limitations, see below) for incidental encoding are met.

Chapter 2 and 6 explored the repetition effect in RSVP. The motivation for these experiments was to test the hypothesis that the memory traces laid by below-threshold (non-reportable) items could facilitate the processing of subsequent instances of the same item. In Chapter 2 we reported the result from two repetition-detection experiments, where instances of the same item were inserted in different RSVP trials. We focussed on whether undetected repetitions increase the probability of detecting a repetition in subsequent presentations of the critical item. That is, we tested whether the process of evidence accumulation can take place even across below-threshold presentations. The results showed that the probability of detecting a repetition did not increase as the number of undetected repetitions increased. This supported the notion that, at least when repetitions are considerably spaced, the memory traces of below-threshold items are too weak to facilitate the processing of subsequent presentations. In conclusion, the evidence reported in Chapter 2 revealed that one of the key properties of episodic representations, their durability, cannot be attributed to the traces left by RSVP items.

Chapter 6 explored further the durability of memory representations in RSVP experiments. In the experiment reported in this chapter, repetitions were inserted in the same stream. Crucially, the spacing between presentations was manipulated from 1 to 3 distractors. In agreement with the evidence in Chapter 2, the accuracy at detecting a repetition when the repetitions were spaced by 3 distractors was consistent with the hypothesis of no (subliminal) evidence accumulation. Interestingly, when the repetitions were separated by 1 and 2 distractors we observed that repetition detection was higher relative to that expected in case of no subliminal evidence accumulation. We concluded that subliminal memory traces can have an impact on subsequent presentations only when the distance between presentations was less than 300 ms. This time frame is what is expected in priming experiments, where the facilitation is naturally accounted for by pre-activation of the same representations. Consequently, we think that the observed facilitatory effect is better accounted for by priming than by the encoding of episodic events.

Chapter 3 investigated recency effects in RSVP in recognition and free-recall. Free-recall is a good test for the recollection mechanism of memory retrieval while recognition is influenced by both recollection and familiarity (Yonelinas, 2002). We explored the recency effect by testing the last 3 items in fast and slow RSVP sequences. The most interesting result was the disproportionately large recency effect observed in fast streams in the free-recall test. Participants reported the last item in the stream in just 20% of the trials, while reportability was very close to 0 for the second to last and the third to last items. This result indicated that even for task-relevant items, reportability is severely impaired. We concluded that while the quality of the memory traces for items in RSVP supported the extraction of familiarity signals, it was not enough to support recollection.

Finally, in Chapter 5 we addressed the list-length effect, and specifically the possibility of proactive interference in RSVP streams. If Chapters 2 and 6 explored a facilitatory effect from previous items in a stream, Chapter 5 investigated interference effects. We created short and long lists of items presented in fast and slow RSVP sequences. The retention interval was kept constant between short and long streams and, consequently, in the long sequences, the critical items were preceded by at least 24 more items than in the short lists. Although the experiment did not reveal an LLE on d', it showed that the difference in RTs between short and long streams was larger in the slow streams relative to the fast streams. Fitting the accuracy and RT data with a DDM model showed an LLE in the threshold parameter for the slow streams but not for the fast streams. Interestingly, the threshold parameter is assumed to vary with the subjective decisions of the participants. This implies that participants adjusted their conservatism in response to the LLE manipulation, but only in the slow streams. Therefore, the presence of 24+ items in RSVP failed to influence the criterion set by the participants in a recognition test in RSVP. This could indicate that the participants estimated that their chance of success was independent of the number of items presented in the fast streams.

2. Limitations and future directions

A central limitation of the present work is that we have not addressed the issue of phenomenal consciousness. All our conclusions are based on the capacity of participants for reporting stimuli presented in RSVP. Although reportability remains the critical criterion to attribute consciousness, there have been some studies reporting dissociations between reports and subjective visibility. In RSVP, Pincham et al. (2016) found low visibility and high reportability for items presented in the lag-1 position of the AB. Interestingly, they found that visibility increased as reportability improved for the other positions in the stream. We suspect that the dissociation between visibility and report might only be observed in the particular case of lag-1 sparing. Therefore, although adding a subjective visibility question to the present experiments may not add much (because we are not looking at lag-1), it might be worth incorporating in future experiments to investigate this issue. It could be interesting for experiments such as the one reported in Chapter 3. We observed good recognition performance for the second-to-last and third-to-last items but a severely impaired free-recall performance for those positions. We think that subjective visibility could be particularly informative for these cases, where there is a very poor recollection and a relatively high familiarity.

One limitation of the present studies concerns the nature of the stimuli employed. We used words as the items to build our RSVP sequences. Although using the same type of stimuli increases the consistency and interpretability of the present studies, it hinders the comparison with some of the most relevant studies reported in the literature. For example, Potter (1976) is a milestone for the rationale of the present work, and they used streams of pictures. Other relevant studies (Endress & Potter, 2014, Thunell & Thorpe, 2019), also used pictures as stimuli for experiments addressing similar issues as this thesis. The main reason for the use of words as stimuli in the present experiments is that some of them required an explicit report in the form of free-recall (Chapter 3) or the report of the repeated item (Chapter 2). We think that for the main contribution of this thesis, the nature of the stimuli is not particularly relevant. The "missing memory imprint", the gap between recognition and detection seems to be robust enough. Perhaps, for more fine-grained effects, such as the evidence accumulation observed in Chapter 6, a different pattern could emerge with pictorial stimuli. If the residual activation trace of pictures is stronger than that of words, then perhaps even with a gap of three distractors, a pattern of evidence accumulation could be observed.

Perhaps the most interesting line of research emerging from this work might be the exploration of the neural correlates of some of the present findings. One common result observed in EEG studies using RSVP is a P3 component associated with the detection of searched-for items (Bowman et al., 2013). In fact, it has been argued that the P3 component might be a neural correlate of consciousness (Lamy et al., 2009). It might be particularly interesting to register EEG activity in an experiment similar to our repetition detection experiments (Chapters 2 and 6). If our interpretation is correct, then the ERPs associated with each presentation before the detection of the repetition should show some interesting differences. Undetected repetitions should not generate a P3. However, when repetitions are close enough (spaced by 1 or 2 distractors), we have concluded that the pattern of evidence accumulation might be related to priming. Therefore, the ERPs of the instances primed by earlier presentations should reveal this facilitation (perhaps a reduced N400).

In addition, collecting EEG data associated with RSVP stimuli could shed some light on a phenomenon that we still cannot account for. It is about the difference between Recollection and Familiarity (Yonelinas, 2002). We explored this issue in Chapter 3. We found that only in a handful of trials participants could report the last item in the stream in a free-recall task. The reportability of the second-to-last item, for example, was much smaller. In contrast, items in these positions were recognized well above chance. It might be potentially interesting to obtain the ERPs for both, recollected and recognized items. Sure, there is the issue that some of the recognition responses could be based on recollective processes (Yonelinas, 2002). So, additional measures (confidence) should be collected to address this issue. Still, the question remains: perhaps, the ERPs commonly associated with conscious processing are to be found more closely linked to to-be-recollected items than to items recognized on the basis of familiarity.

Another promising new direction related to the present work is the investigation of metacognitive processes associated with recognition responses. In the experiment reported in Chapter 6, the repetition detection task, we observed an improvement in sensitivity in the recognition task as the number of repetitions increased. We think that calculating the meta-d' (Maniscalco & Lau, 2012) for the recognition task could shed some light on to what extent the observed sensitivity is based on memory traces accessible to introspection. In the repetition detection task, we interpreted that there was some unconscious evidence accumulation for repetitions spaced by 1 or 2 distractors. We hypothesize that this unconscious facilitation may affect the d' for recognition but could not influence the meta-d'. Metacognitive efficiency (Fleming & Lau, 2014) might increase with the spacing of the repetitions, as the unconscious evidence accumulation effect vanishes.

Finally, one important test of the Tokenized Percept Hypothesis would be to explore the relationship between unconscious perception and order information, one of the defining features of episodic processing. In this thesis, we have investigated the idea of the unconscious being "episodically blind" by testing the repetition effect and the overall reportability of RSVP items. However, the encoding of order information is as crucial to our main hypothesis as the capacity for detecting repetitions. Although there is no obvious way of testing the order of events that are barely reportable, an experiment showing that order information is lost for items in RSVP would support the central tenets of the Tokenized Percept Hypothesis.

3. Implications and conclusion

The main contribution of the present work is to offer a comprehensive exploration of the episodic encoding of items in RSVP. The "repetition detection" chapters, 2 and 6, offered an alternative interpretation and complementary evidence to the findings of Endress and Potter (2014) and Thunell and Thorpe (2019). We showed in Chapter 2 that the improvement in recognition performance for repeated items observed in Endress and Potter (2014) is, most likely, based on conscious processes. This is so because we rejected the hypothesis of evidence accumulation for subliminal stimuli. Alongside the evidence from Chapter 6, we have "parametrized" the spacing conditions in which unconscious evidence accumulation could take place. We concluded that repetitions spaced for 300 ms or more do not benefit from earlier subliminal instances of the same stimulus. And the closer the presentations, the greater the facilitation. Furthermore, this contradicts the classic "spacing effect" where the benefits of repetition tend to be more easily observed as the distance between presentations increases (Raaijmakers, 2003; Verkoeijen et al., 2005). This seems not to be the case for RSVP.

The experiment in Chapter 4, the "surprise question experiment", with its limitations, is as far as we know the experiment that more directly addressed the reportability of a random distractor in RSVP. The incidental encoding conditions and the surprise question procedure made it possible to test the memory for irrelevant items -arguably, also testing our statistical ingenuity. Although perhaps more evidence is needed to build a more compelling case, our conclusion is that the memory traces of task-irrelevant items do not modulate the performance in an explicit memory test.

Chapter 5 also presented some innovations that could contribute to the field. The registration of reaction times and the drift-diffusion modelling allowed us to make inferences about the subjective part of the recognition task. Interestingly, we found that in RSVP, this subjective process was impervious to the number of items presented. This result complements the findings from the rest of the empirical chapters. It is not only that items in RSVP are barely reportable, but also participants seem to care little about how the presentation of a massive quantity of items (36 in the long lists) might affect their performance. A possible interpretation of this is that many of these items are not consciously perceived, or at the least, leave durable memory traces.

Taken together, these findings reinforce the view that the low recognition performance for items in RSVP reported by Potter (1976) is not simply a matter of forgetting, since even the reportability of the last item in the stream is severely impaired (Chapters 3 and 4). The usual argument of stimuli being consciously perceived but quickly forgotten is challenged by these findings. We think that the most natural interpretation is that many of the stimuli never reach the reportability threshold. Our Tokenized Percept Hypothesis states that these non-reportable items should not leave any episodic trace. The present findings agree with this statement. If it does not facilitate as an episodic representation, does not interfere as an episodic representation, and does not decay as an episodic representation, probably is not an episodic representation.

We consider that one of the main contributions of this thesis is to add the "missing memory imprint" (Potter, 1976) to the range of effects that might offer a clue about the functional correlates of consciousness. However, this effect, per se, is mute about the nature of conscious processes. Kanwisher's (2001) Token Individuation Hypothesis and our Tokenized Percept Hypothesis offer the theoretical tools to establish the relevance of that effect. The tools are Types, Tokens, and a 2-stage process of temporal attention. To put it simply, detection can be achieved on the basis of categorical and/or featural information, Types. Only the "catness" of a cat is needed to detect a cat, but more is needed to declare that detection. Factuality must be explicitly represented (Dienes & Perner, 1999). Factuality and other episodic properties, such as order information, are assigned to first-order representations, Types, in the second stage of processing. By linking Types to higher-order content-independent representations, Token, information that was *in* the system mutates to be information *for* the system. Implicit information about the relationship between representations, such as order information, becomes explicitly represented by the tokenization process. In our view, conscious contents are not only about the immediate percept but also about its episodic context.

The Tokenized Percept Hypothesis has a long way to go before it can be a proper theory of consciousness. Our aim in this thesis was to test one of its core ideas: that the explicit encoding of episodic information is a functional correlate of consciousness. Although the present findings cannot be considered conclusive proof, they certainly provide strong evidence in favour of this claim.

References

Agus, T. R., Thorpe, S. J., & Pressnitzer, D. (2010). Rapid Formation of Robust Auditory Memories: Insights from Noise. *Neuron*, 66(4), 610–618. https://doi.org/10.1016/J.NEURON.2010.04.014

Albrecht, T., & Vorberg, D. (2010). Long-lasting effects of briefly flashed words and pseudowords in ultrarapid serial visual presentation. *Journal of Experimental Psychology*. Learning, Memory, and Cognition, 36(5), 1339–1345. doi:10.1037/a0019999

Anderson, A. K., & Phelps, E. A. (2001). Lesions of the human amygdala impair enhanced perception of emotionally salient events. *Nature* 2001 411:6835, 411(6835), 305–309. https://doi.org/10.1038/35077083

Atkinson, R. C., & Shiffrin, R. M. (1971). The Control of Short-Term Memory. *Scientific American*, 225(2), 82–91. http://www.jstor.org/stable/24922803

Avilés, A., Bowman, H., & Wyble, B. (2020). On the limits of evidence accumulation of the preconscious percept. *Cognition*, 195, 104080. https://doi.org/10.1016/j.cognition.2019.104080

Baars, B. J. (1993). A cognitive theory of consciousness. Cambridge University Press.

Baars, B. J. (2002). The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Sciences*, 6(1), 47–52. doi:10.1016/S1364-6613(00)01819-2

Baars, B. J., & Franklin, S. (2003). How conscious experience and working memory interact. *Trends in Cognitive Sciences*, 7(4), 166–172. https://doi.org/10.1016/S1364-6613(03)00056-1

Baars, B. J. (2017). The Global Workspace Theory of Consciousness. In The Blackwell Companion to Consciousness (pp. 227–242). John Wiley & Sons, Ltd. https://doi.org/https://doi.org/10.1002/9781119132363.ch16

Baddeley, A. D., & Hitch, G. (1974). Working memory. In G.H. Bower (Ed.), The psychology of learning and motivation: Advances in research and theory (Vol. 8, pp. 47–89). New York: Academic Press.

Baddeley, A. (2000). The episodic buffer: a new component of working memory? *Trends in Cognitive Sciences*, 4(11), 417–423. https://doi.org/10.1016/S1364-6613(00)01538-2

Baddeley, A. (2003). Working memory: Looking back and looking forward. *Nature Reviews Neuroscience*, 4(10), 829–839. https://doi.org/10.1038/nrn1201

Balota, D. A., Yap, M. J., Cortese, M. J., Hutchison, K. A., Kessler, B., Loftis, B., … Treiman, R. (2007). The english lexicon project. *Behavior Research Methods*, 39(3), 445–459. doi:10.3758/BF03193014

Barnard, P. J., Scott, S., Taylor, J., May, J., & Knightley, W. (2004). Paying attention to meaning. *Psychological Science*, 15(3), 179-186.

Barnard, P., & Bowman, H. (2004). Rendering Information Processing Models of Cognition and Affect Computationally Explicit: Distributed Executive Control and the Deployment of Attention. *Cognitive Science Quarterly*, 3(3), 297–328.

Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi:10.18637/jss.v067.i01

Bayne, T. (2018). On the axiomatic foundations of the integrated information theory of consciousness. *Neuroscience of Consciousness*, 2018(1). https://doi.org/10.1093/NC/NIY007

Bergström, F., & Eriksson, J. (2014). Maintenance of non-consciously presented information engages the prefrontal cortex. *Frontiers in Human Neuroscience*, 8, 938. doi:10.3389/fnhum.2014.00938

Block, N. (1995). On a confusion about a function of consciousness. *Behavioral and Brain Sciences*, 18(2), 227–247. https://doi.org/10.1017/S0140525X00038188

Block, N. (2007). Consciousness, accessibility, and the mesh between psychology and neuroscience. *Behavioral and Brain Sciences*, 30(5–6), 481–99; discussion 499. doi:10.1017/S0140525X07002786

Bowman, H., & Avilés, A. (2021, March 31). Fragile Memories for Fleeting Percepts. https://doi.org/10.31234/osf.io/fr3bg

Bowman, H., & Avilés, A. (2022). No Subliminal Memory for Spaced Repeated Images in Rapid-Serial-Visual-Presentation Streams. *Psychological Science*. https://doi.org/10.1177/09567976221100229

Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological Review*, 114(1), 38–70. https://doi.org/10.1037/0033-295X.114.1.38

Bowman, H., Filetti, M., Alsufyani, A., Janssen, D., & Su, L. (2014). Countering countermeasures: Detecting identity lies by detecting conscious breakthrough. *PLoS ONE*, 9(3). https://doi.org/10.1371/journal.pone.0090595

Bowman, H., Filetti, M., Janssen, D., Su, L., Alsufyani, A., & Wyble, B. (2013). Subliminal Salience Search Illustrated: EEG Identity and Deception Detection on the Fringe of Awareness. *PLoS ONE*, 8(1), e54258. https://doi.org/10.1371/journal.pone.0054258

Bowman, H., Wyble, B., Chennu, S., & Craston, P. (2008). A reciprocal relationship between bottom-up trace strength and the attentional blink bottleneck: Relating the LC-NE and ST2 models. *Brain Research*, 1202, 25–42. https://doi.org/10.1016/j.brainres.2007.06.035

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision*, 10(4), 433–436. https://doi.org/10.1163/156856897X00357

Broadbent, D. E., & Broadbent, M. H. P. (1987). From detection to identification: Response to multiple targets in rapid serial visual presentation. *Perception & Psychophysics* 1987 42:2, 42(2), 105–113. https://doi.org/10.3758/BF03210498

Brown, R., Lau, H., & LeDoux, J. E. (2019). Understanding the Higher-Order Approach to Consciousness. *Trends in Cognitive Sciences*, 23(9), 754–768. https://doi.org/10.1016/J.TICS.2019.06.009

Capretto, T., Piho, C., Kumar, R., Westfall, J., Yarkoni, T., & Martin, O. A. (2022). Bambi: A Simple Interface for Fitting Bayesian Linear Models in Python. *Journal of Statistical Software*, 103(15), 1–29. https://doi.org/10.18637/JSS.V103.I15

Carrasco, M., Ling, S., & Read, S. (2004). Attention alters appearance. *Nature neuroscience*, *7*(3), 308-313.

Carruthers, P., & Gennaro, R. (2020). Higher-Order Theories of Consciousness. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Fall 2020). Metaphysics Research Lab, Stanford University.

Cary, M., & Reder, L. M. (2003). A dual-process account of the list-length and strength-based mirror effects in recognition. *Journal of Memory and Language*, 49(2), 231–248. https://doi.org/10.1016/S0749-596X(03)00061-5

Cepeda, N. J., Pashler, H., Vul, E., Wixted, J. T., & Rohrer, D. (2006). Distributed practice in verbal recall tasks: A review and quantitative synthesis. *Psychological Bulletin*, 132(3), 354–380. https://doi.org/10.1037/0033-2909.132.3.354

Chalmers, D. J. (1996). The Conscious Mind: In Search of a Fundamental Theory. Oxford University Press.

Chen, H., & Wyble, B. (2016). Attribute amnesia reflects a lack of memory consolidation for attended information. *Journal of Experimental Psychology: Human Perception and Performance*, 42(2), 225–234. https://doi.org/10.1037/xhp0000133

Chong, T. T. J., Husain, M., & Rosenthal, C. R. (2014). Recognizing the unconscious. *Current Biology*, 24(21), R1033–R1035. https://doi.org/10.1016/j.cub.2014.09.035

Chun, M. M., & Potter, M. (1995). A Two-Stage Model for Multiple Target Detection in Rapid Serial Visual Presentation. *Journal of Experimental Psychology: Human Perception and Performance*, 21(1), 109–127. https://doi.org/10.1037/0096-1523.21.1.109

Cleeremans, A. (2011). The radical plasticity thesis: How the brain learns to be conscious. *Frontiers in Psychology*, 2(MAY), 1–12. https://doi.org/10.3389/fpsyg.2011.00086

Cleeremans, A., & Tallon-Baudry, C. (2022). Consciousness matters: phenomenal experience has functional value. *Neuroscience of Consciousness*, 2022(1), 1–11. https://doi.org/10.1093/NC/NIAC007

Cohen, M. A., & Chun, M. M. (2017a). Studying consciousness through inattentional blindness, change blindness, and the attentional blink. In The Blackwell companion to consciousness (pp. 539–550). Wiley Online Library.

Cohen, M. A., & Dennett, D. C. (2011). Consciousness cannot be separated from function. *Trends in Cognitive Sciences*, 15(8), 358–364. https://doi.org/10.1016/j.tics.2011.06.008

Cohen, N. J., Poldrack, R. A., & Eichenbaum, H. (1997). Memory for Items and Memory for Relations in the Procedural/Declarative Memory Framework. *Memory*, 5(1–2), 131–178. https://doi.org/10.1080/741941149

Collins, A. M., & Loftus, E. F. (1975). A spreading-activation theory of semantic processing. *Psychological Review*, 82(6), 407–428. https://doi.org/10.1037/0033-295X.82.6.407

Cowan, N. (2010). The Magical Mystery Four: How is Working Memory Capacity Limited, and Why? *Current Directions in Psychological Science*, 19(1), 51. https://doi.org/10.1177/0963721409359277

Cowan, N. (2016). Working memory capacity: Classic edition. In Working Memory Capacity: Classic Edition. Routledge. https://doi.org/10.4324/9781315625560

Craston, P., Wyble, B., Chennu, S., & Bowman, H. (2009). The attentional blink reveals serial working memory encoding: Evidence from virtual and human event-related potentials. *Journal of cognitive neuroscience*, 21(3), 550-566.

DeCarlo, L. T. (1998). Signal Detection Theory and Generalized Linear Models. *Psychological Methods*, 3(2), 186–205. https://doi.org/10.1037/1082-989X.3.2.186

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

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

Dehaene, S., Sergent, C., & Changeux, J. P. (2003). A neuronal network model linking subjective reports and objective physiological data during conscious perception. *Proceedings of the National Academy of Sciences*, 100(14), 8520–8525. https://doi.org/10.1073/PNAS.1332574100

Dehaene, S. (2014). *Consciousness and the brain: Deciphering how the brain codes our thoughts*. Penguin.

Dehaene, S., Changeux, J. P., Naccache, L., Sackur, J., & Sergent, C. (2006). Conscious, preconscious, and subliminal processing: a testable taxonomy. *Trends in cognitive sciences*, *10*(5), 204-211.

Dehaene, S., & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: basic evidence and a workspace framework. *Cognition*, *79*(1-2), 1-37.

Dennett, D. C. (1993). Consciousness explained. Penguin uk.

Dennis, S., & Humphreys, M. S. (2001). A context noise model of episodic word recognition. *Psychological Review*, 108(2), 452–478. https://doi.org/10.1037/0033-295X.108.2.452

Desimone, R., & Duncan, J. (1995). Neural Mechanism of Selective Visual Attention. *Annual Review of Neuroscience*, 18, 193-222.

Dienes, Z., & Perner, J. (1999). A theory of implicit and explicit knowledge. *Behavioral and Brain Sciences*, 22(5), 735–808. https://doi.org/10.1017/S0140525X99002186

Droege, P. (2009). Now or never: How consciousness represents time. *Consciousness and Cognition*, 18(1), 78–90. https://doi.org/10.1016/j.concog.2008.10.006

Duss, S. B., Reber, T. P., Hänggi, J., Schwab, S., Wiest, R., Müri, R. M., Brugger, P., Gutbrod, K., & Henke, K. (2014). Unconscious relational encoding depends on hippocampus. *Brain*, 137(12), 3355–3370. https://doi.org/10.1093/brain/awu270

Dux, P. E., Coltheart, V., & Harris, I. M. (2006). On the fate of distractor stimuli in rapid serial visual presentation. *Cognition*, 99(3), 355–382. https://doi.org/10.1016/j.cognition.2005.04.002

Ebbinghaus, H. (1964). Memory: A contribution to experimental psychology. NewYork: Dover Publications.

Endress, A. D., & Potter, M. (2014). Something from (almost) nothing: buildup of object memory from forgettable single fixations. *Attention, Perception, and Psychophysics*, 76(8), 2413–2423. https://doi.org/10.3758/s13414-014-0706-3

Faivre, N., Mudrik, L., Schwartz, N., & Koch, C. (2014). Multisensory Integration in Complete Unawareness: Evidence From Audiovisual Congruency Priming. *Psychological Science*, 25(11), 2006–2016.

Ferrand, L. (1996). The masked repetition priming effect dissipates when increasing the inter-stimulus interval: Evidence from word naming. *Acta Psychologica*, 91(1), 15–25. doi:10.1016/0001-6918(95)00010-0

Fleming, S. M. (2020). Awareness as inference in a higher-order state space. *Neuroscience of Consciousness*, 2020(1). https://doi.org/10.1093/NC/NIZ020

Fleming, S. M., & Lau, H. C. (2014). How to measure metacognition. *Frontiers in Human Neuroscience*, 8(JULY), 443. https://doi.org/10.3389/FNHUM.2014.00443/ABSTRACT

Fodor, J. (2004). You can't argue with a novel. London Review of Books, 26(5). https://www.lrb.co.uk/the-paper/v26/n05/jerry-fodor/you-can-t-argue-with-a-novel

Forster, K. (1970). Visual perception of rapidly presented text segments of varying complexity. *Perception & Psychophysics*, 8, 215–221.

Frith, C. D., & Metzinger, T. (2016). What's the Use of Consciousness? The Pragmatic Turn, 193–214. https://doi.org/10.7551/mitpress/9780262034326.003.0012

Graf, P., & Schacter, D. L. (1985). Implicit and Explicit Memory for New Associations in Normal and Amnesic Subjects. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 11(3), 501–518. https://doi.org/10.1037/0278-7393.11.3.501

Hassin, R. R. (2013). Yes It Can: On the Functional Abilities of the Human Unconscious. *Perspectives on Psychological Science*, 8(2), 195–207. https://doi.org/10.1177/1745691612460684

Henke, K. (2010). A model for memory systems based on processing modes rather than consciousness. *Nature Reviews. Neuroscience*, 11(7), 523–532. https://doi.org/10.1038/NRN2850

Hintzman, D. L. (1986). "Schema Abstraction" in a Multiple-Trace Memory Model. *Psychological Review*, 93(4), 411–428. https://doi.org/10.1037/0033-295X.93.4.411

Hoffman, M. D., & Gelman, A. (2014). The No-U-Turn Sampler: Adaptively Setting Path Lengths in Hamiltonian Monte Carlo. In *Journal of Machine Learning Research* (Vol. 15). http://mcmc-jags.sourceforge.net

Huxley, T. H., 1874, "On the Hypothesis that Animals are Automata, and its History", The Fortnightly Review 16 (New Series): 555–580. Reprinted in Method and Results: Essays by Thomas H. Huxley, New York: D. Appleton and Company, 1898.

Intraub, H. (1981). Rapid conceptual identification of sequentially presented pictures. *Journal of Experimental Psychology: Human Perception and Performance*, 7(3), 604.

Jiang, Y., Costello, P., Fang, F., Huang, M., & He, S. (2006). A gender- and sexual orientation-dependent spatial attentional effect of invisible images. *Proceedings of the National Academy of Sciences*, 103(45), 17048–17052. https://doi.org/10.1073/PNAS.0605678103

Johnson, M. K., & Raye, C. L. (1981). Reality monitoring. *Psychological Review*, 88(1), 67–85. https://doi.org/10.1037/0033-295X.88.1.67

Jones, W., Pincham, H., Gootjes-Dreesbach, E. L., & Bowman, H. (2020). Fleeting Perceptual Experience and the Possibility of Recalling Without Seeing. *Scientific Reports*, 10(1). https://doi.org/10.1038/S41598-020-64843-2

Kahana, M. J., Diamond, N. B., & Aka, A. (2022, April 6). Laws of Human Memory. https://doi.org/10.31234/osf.io/aczu

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.

Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79(1–2), 89–113. https://doi.org/10.1016/S0010-0277(00)00125-6

Kanwisher, N. G. (1987). Repetition blindness: Type recognition without token individuation. *Cognition*, 27(2), 117–143. https://doi.org/10.1016/0010-0277(87)90016-3

Kentridge, R. W., Heywood, C. A., & Weiskrantz, L. (1999). Attention without awareness in blindsight. *Proceedings of the Royal Society of London. Series B: Biological Sciences*, *266*(1430), 1805-1811.

Kinnell, A., & Dennis, S. (2011). The list length effect in recognition memory: An analysis of potential confounds. *Memory and Cognition*, 39(2), 348–363. https://doi.org/10.3758/s13421-010-0007-6

Kleiner, M. (2010). Visual stimulus timing precision in Psychtoolbox-3: Tests, pitfalls and solutions. Perception 39 ECVP Abstract Supplement, 189. https://doi.org/10.1068/v100546

Kruschke, J. K. (2013). Bayesian estimation supersedes the t test. *Journal of Experimental Psychology: General*, 142(2), 573–603. https://doi.org/10.1037/a0029146

Lachter, J., Forster, K. I., & Ruthruff, E. (2004). Forty-five years after Broadbent (1958): still no identification without attention. *Psychological review*, 111(4), 880.

Lakens, D., & Caldwell, A. R. (2021). Simulation-Based Power Analysis for Factorial Analysis of Variance Designs. *Advances in Methods and Practices in Psychological Science*, 4(1). https://doi.org/10.1177/2515245920951503

Lamme, V. A. F. (2003). Why visual attention and awareness are different. *Trends in Cognitive Sciences*, 7(1), 12–18. https://doi.org/10.1016/S1364-6613(02)00013-X

Lamme, V. A. F. (2010). How neuroscience will change our view on consciousness. *Cognitive Neuroscience*, 1(3), 204–220. https://doi.org/10.1080/17588921003731586

Lamme, V. A., Zipser, K., & Spekreijse, H. (1998). Figure-ground activity in primary visual cortex is suppressed by anesthesia. *Proceedings of the National Academy of Sciences*, *95*(6), 3263-3268.

Lamme V. A. F. (2018). Challenges for theories of consciousness: seeing or knowing, the missing ingredient and how to deal with panpsychism. *Philosophical transactions of the Royal Society of London. Series B, Biological sciences*, *373*(1755), 20170344. https://doi.org/10.1098/rstb.2017.0344

Lamy, D., Salti, M., & Bar-Haim, Y. (2009). Neural Correlates of Subjective Awareness and Unconscious Processing: An ERP Study. *Journal of Cognitive Neuroscience*, 21(7), 1435–1446. https://doi.org/10.1162/jocn.2009.21064

Lau, H. C., & Passingham, R. E. (2007). Unconscious activation of the cognitive control system in the human prefrontal cortex. *Journal of Neuroscience*, 27(21), 5805–5811. https://doi.org/10.1523/JNEUROSCI.4335-06.2007

Lau, H. (2019, June 10). Consciousness, Metacognition, & Perceptual Reality Monitoring. https://doi.org/10.31234/osf.io/ckbyf

Lau, H. (2022). In Consciousness we Trust: The Cognitive Neuroscience of Subjective Experience. Oxford University Press.

Mack, A., & Rock, I. (1998). Inattentional blindness: Perception without attention. Visual Attention, 8, 55–76.

Maniscalco, B., & Lau, H. (2012). A signal detection theoretic approach for estimating metacognitive sensitivity from confidence ratings. *Consciousness and Cognition*, 21(1), 422–430. https://doi.org/10.1016/j.concog.2011.09.021

Martini, P., & Maljkovic, V. (2009). Short-term memory for pictures seen once or twice. *Vision Research*, 49(13), 1657–1667. https://doi.org/10.1016/J.VISRES.2009.04.007

McClelland, J. L. (1979). On the time relations of mental processes: an examination of systems of processes in cascade. *Psychological review*, 86(4), 287.

McClelland, J. L., & Rumelhart, D. E. (1981). An interactive activation model of context effects in letter perception: I. An account of basic findings. *Psychological Review*, 88(5), 375–407. https://doi.org/10.1037/0033-295X.88.5.375

Morys-Carter, W. L. (2021, May 18). ScreenScale [Computer software].

Most, S. B., Chun, M. M., Widders, D. M., & Zald, D. H. (2005). Attentional rubbernecking: Cognitive control and personality in emotion-induced blindness. *Psychonomic bulletin & review*, 12(4), 654-661.

Müller, G. E., & Pilzecker, A. (1900). Experimentelle beiträge zur lehre vom gedächtniss (Vol. 1). JA Barth.

Naccache, L., Blandin, E., & Dehaene, S. (2002). Unconscious masked priming depends on temporal attention. *Psychological Science*, 13(5), 416–424. https://doi.org/10.1111/1467-9280.00474

Neisser, U. (2014). Cognitive psychology: Classic edition. Psychology press.

Nieuwenstein, M. R., & Potter, M. (2006). Temporal limits of selection and memory encoding: A comparison of whole versus partial report in rapid serial visual presentation. *Psychological Science*, 17(6), 471–475. https://doi.org/10.1111/j.1467-9280.2006.01730.x

Nieuwenstein, M., Burg, E. van der, & Potter, M. (2009).

Norman, D.A., Shallice, T. (1986). Attention to Action. In: Davidson, R.J., Schwartz, G.E., Shapiro, D. (eds) Consciousness and Self-Regulation. Springer, Boston, MA. https://doi.org/10.1007/978-1-4757-0629-1_1

Ohrt, D. D., & Gronlund, S. D. (1999). List-length effect and continuous memory: Confounds and solutions. In C. Izawa (Ed.), On human memory: Evolution, progress, and reflections on the 30th anniversary of the Atkinson-Shiffrin model (pp. 105–125). Lawrence Erlbaum Associates Publishers

Paulewicz, B., & Blaut, A. (2020). The bhsdtr package: a general-purpose method of Bayesian inference for signal detection theory models. *Behavior Research Methods*, 52(5), 2122–2141. https://doi.org/10.3758/S13428-020-01370-Y/FIGURES/10

Peirce, J., Gray, J. R., Simpson, S., MacAskill, M., Höchenberger, R., Sogo, H., Kastman, E., & Lindeløv, J. K. (2019). PsychoPy2: Experiments in behavior made easy. *Behavior Research Methods*, 51(1), 195–203. https://doi.org/10.3758/s13428-018-01193-y

Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, 10(4), 437–442. https://doi.org/10.1163/156856897X00366

Pincham, H. L., Bowman, H., & Szucs, D. (2016). The experiential blink: Mapping the cost of working memory encoding onto conscious perception in the attentional blink. *Cortex*, 81, 35–49. https://doi.org/10.1016/J.CORTEX.2016.04.007

Postle, B. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*(1), 23-38. https://doi.org/10.1016/j.neuroscience.2005.06.005

Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of experimental psychology: human learning and memory*, 2(5), 509.

Potter, M. C., Chun, M. M., Banks, B. S., & Muckenhoupt, M. (1998). Two attentional deficits in serial target search: the visual attentional blink and an amodal task-switch deficit. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 24(4), 979.

Potter, M. C., Staub, A., & O'Connor, D. H. (2002). The time course of competition for attention: attention is initially labile. *Journal of Experimental Psychology: Human Perception and Performance*, 28(5), 1149.

Potter, M., Wyble, B., Hagmann, C. E., & McCourt, E. S. (2014). Detecting meaning in RSVP at 13 ms per picture. *Attention, Perception, and Psychophysics*, 76(2), 270–279. https://doi.org/10.3758/s13414-013-0605-z

Raaijmakers, J. G. W. (2003). Spacing and repetition effects in human memory: Application of the SAM model. *Cognitive Science*, 27(3), 431–452. https://doi.org/10.1016/S0364-0213(03)00007-7

Ratcliff, R. (1978). A theory of memory retrieval. *Psychological Review*, 85(2), 59–108. https://doi.org/10.1037/0033-295X.85.2.59

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. https://doi.org/10.1037//0096-1523.18.3.849

Reber, T. P., & Henke, K. (2011). Rapid formation and flexible expression of memories of subliminal word pairs. *Frontiers in Psychology*, 2(NOV), 1–11. https://doi.org/10.3389/fpsyg.2011.00343

Reder, L. M., Nhouyvanisvong, A., Schunn, C. D., Ayers, M. S., Angstadt, P., & Hiraki, K. (2000). A mechanistic account of the mirror effect for word frequency: a computational model of remember-know judgments in a continuous recognition paradigm. *Journal of Experimental Psychology. Learning, Memory, and Cognition*, 26(2), 294–320. doi:10.1037/0278-7393.26.2.294

Rees, G., Russell, C., Frith, C. D., & Driver, J. (1999). Inattentional Blindness Versus Inattentional Amnesia for Fixated But Ignored Words. *Science*, 286(5449), 2504–2507. https://doi.org/10.1126/science.286.5449.2504

Rensink, R. A., O'Regan, J. K., & Clark, J. J. (1997). To See or not to See: The Need for Attention to Perceive Changes in Scenes. *Psychological Science*, 8(5), 368–373. https://doi.org/10.1111/j.1467-9280.1997.tb00427.x

Robinson, William, "Epiphenomenalism", The Stanford Encyclopedia of Philosophy (Summer 2019 Edition), Edward N. Zalta (ed.), URL = <https://plato.stanford.edu/archives/sum2019/entries/epiphenomenalism/>.

Rosenthal, D. M. (1997). A Theory of Consciousness. In N. Block, O. J. Flanagan, & G. Guzeldere (Eds.), The Nature of Consciousness. MIT Press.

Rosenthal, D. M. (2008). Consciousness and its function. *Neuropsychologia*, 46(3), 829–840. https://doi.org/10.1016/J.NEUROPSYCHOLOGIA.2007.11.012

Schouten, J. F., & Bekker, J. A. M. (1967). Reaction time and accuracy. *Acta Psychologica*, 27(C), 143–153. https://doi.org/10.1016/0001-6918(67)90054-6

Sergent, C., Baillet, S., & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature neuroscience*, 8(10), 1391-1400.

Shallice, T. (1972). Dual functions of consciousness. *Psychological Review*, 79(5), 383–393. https://doi.org/10.1037/H0033135

Shapiro, K. L., Caldwell, J., & Sorensen, R. E. (1997). Personal Names and the Attentional Blink: A Visual "Cocktail Party" Effect. *Journal of Experimental Psychology: Human Perception and Performance*, 23(2), 504–514. https://doi.org/10.1037/0096-1523.23.2.504

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(2), 95–100. https://doi.org/10.1111/j.1467-9280.1997.tb00689.x

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

Sligte, I. G., Scholte, H. S., & Lamme, V. A. (2008). Are there multiple visual short-term memory stores?. *PLOS one*, 3(2).

Soto, D., Mäntylä, T., & Silvanto, J. (2011). Working memory without consciousness. *Current Biology*, 21(22), 49–51. https://doi.org/10.1016/j.cub.2011.09.049

Sperling, G. (1960). The information available in brief visual presentations. *Psychological monographs: General and applied*, 74(11), 1.

Squire, L. R., & Zola, S. M. (1996). Structure and function of declarative and nondeclarative memory systems. *Proceedings of the National Academy of Sciences of the United States of America*, 93(24), 13515–13522.

Sternberg, S. (1975). Memory Scanning: New Findings and Current Controversies. *Quarterly Journal of Experimental Psychology*, 27(1), 1–32. https://doi.org/10.1080/14640747508400459

Stokes, M. G. (2015). 'Activity-silent'working memory in prefrontal cortex: a dynamic coding framework. *Trends in Cognitive Sciences*, *19*(7), 394-405.

Stretch, V., & Wixted, J. T. (1998). Decision rules for recognition memory confidence judgments. *Journal of Experimental Psychology: Learning Memory and Cognition*, 24(6), 1397–1410. https://doi.org/10.1037/0278-7393.24.6.1397

Strong, E. K. (1912). The effect of length of series upon recognition memory. *Psychological Review*, 19(6), 447–462. https://doi.org/10.1037/H0069812

Su, L., Bowman, H., & Barnard, P. (2011). Glancing and then looking: On the role of body, affect, and meaning in cognitive control. *Frontiers in Psychology*, 2(DEC), 1–23. https://doi.org/10.3389/fpsyg.2011.00348

Subramaniam, S., Biederman, I., & Madigan, S. (2000). Accurate identification but no priming and chance recognition memory for pictures in RSVP sequences. *Visual Cognition*, 7(4), 511–535. doi:10.1080/135062800394630

Thunell, E., & Thorpe, S. J. (2019). Memory for Repeated Images in Rapid-Serial-Visual-Presentation Streams of Thousands of Images. *Psychological Science*, 30(7), 989–1000. https://doi.org/10.1177/0956797619842251

Tononi, G. (2017a). Integrated Information Theory of Consciousness. In The Blackwell Companion to Consciousness (pp. 621–633). John Wiley & Sons, Ltd. https://doi.org/https://doi.org/10.1002/9781119132363.ch44

Tononi, G. (2017b). The Integrated Information Theory of Consciousness. In The Blackwell Companion to Consciousness (pp. 243–256). John Wiley & Sons, Ltd. https://doi.org/https://doi.org/10.1002/9781119132363.ch17

Tononi, G., & Koch, C. (2008). The Neural Correlates of Consciousness. *Annals of the New York Academy of Sciences*, 1124(1), 239–261. https://doi.org/10.1196/ANNALS.1440.004

Townsend, J. T., & Ashby, F. G. (1983). The stochastic modelling of elementary psychological processes Cambridge University Press. Cambridge, UK.

Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136. https://doi.org/10.1016/0010-0285(80)90005-5

Trübutschek, D., Marti, S., Ojeda, A., King, J. R., Mi, Y., Tsodyks, M., & Dehaene, S. (2017). A theory of working memory without consciousness or sustained activity. *ELife*, 6, 1–46. https://doi.org/10.7554/eLife.23871

Tsuchiya, N., & Koch, C. (2005). Continuous flash suppression reduces negative afterimages. *Nature Neuroscience*, 8(8), 1096–1101. https://doi.org/10.1038/nn1500

Tulving, E. (1985). Memory and consciousness. Canadian Psychology / Psychologie Canadienne, 26(1), 1–12. https://doi.org/10.1037/H0080017

Tulving, E. (2001). Episodic memory and common sense: how far apart. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, 356(1413), 1505–1515. https://doi.org/10.1098/RSTB.2001.0937

Underwood, B. J. (1957). Interference and forgetting. *Psychological Review*, 64(1), 49–60. https://doi.org/10.1037/H0044616

van Gaal, S., Ridderinkhof, K. R., Scholte, H. S., & Lamme, V. A. F. (2010). Unconscious activation of the prefrontal no-go network. *Journal of Neuroscience*, 30(11), 4143–4150. https://doi.org/10.1523/JNEUROSCI.2992-09.2010

Verkoeijen, P. P. J. L., Rikers, R. M. J. P., & Schmidt, H. G. (2005). Limitations to the spacing effect: demonstration of an inverted u-shaped relationship between repetition spacing and free recall. *Experimental Psychology*, 52(4), 257–263. https://doi.org/10.1027/1618-3169.52.4.257

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(6), 1656.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance*, 32(6), 1436–1451. https://doi.org/10.1037/0096-1523.32.6.1436

Voss, A., Rothermund, K., & Voss, J. (2004). Interpreting the parameters of the diffusion model: An empirical validation. *Memory and Cognition*, 32(7), 1206–1220. https://doi.org/10.3758/BF03196893

Wagenmakers, E.-J. (2007). A practical solution to the pervasive problems of p values. *Psychonomic Bulletin & Review*, 14(5), 779–804. doi:10.3758/BF03194105

Warriner, A. B., Kuperman, V., & Brysbaert, M. (2013). Norms of valence, arousal, and dominance for 13,915 English lemmas. *Behavior Research Methods*, 45(4), 1191–1207. doi:10.3758/s13428-012-0314-x

Wetzel, L. (2018). Types and Tokens. In E. N. Zalta (Ed.), The Stanford Encyclopedia of Philosophy (Fall 2018). Metaphysics Research Lab, Stanford University.

Wetzels, R., Wetzels, R., Raaijmakers, J. G. W., Jakab, E., & Wagenmakers, E.-J. (2009). How to quantify support for and against the null hypothesis: A flexible WinBUGS implementation of a default Bayesian t test. *Psychonomic Bulletin & Review*, 16(4), 752–760. https://doi.org/10.3758/PBR.16.4.752

Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological Review*, 114(1), 152–176. https://doi.org/10.1037/0033-295X.114.1.152

Wyble, B., Bowman, H., & Nieuwenstein, M. (2009). The Attentional Blink Provides Episodic Distinctiveness: Sparing at a Cost. *Journal of Experimental Psychology: Human Perception and Performance*, 35(3), 787–807. https://doi.org/10.1037/a0013902

Wyble, B., Wade, G., & Hess, M. (2018, January 31). Stream. Retrieved from osf.io/tdvxm

Yonelinas, A. P. (2002). The nature of recollection and familiarity: A review of 30 years of research. *Journal of Memory and Language*, 46(3), 441–517. https://doi.org/10.1006/jmla.2002.2864

Zhang, J., & Rowe, J. B. (2014). Dissociable mechanisms of speed-accuracy tradeoff during visual perceptual learning are revealed by a hierarchical drift-diffusion model. *Frontiers in Neuroscience*, 0(8 APR), 69. https://doi.org/10.3389/FNINS.2014.00069/BIBTEX

Appendices

Appendix 1

Appendix for Chapter 2

Appendix 1.1

Mathematical characterisation

Assume a sequence of N trials in which an item repeats, index these trials with the natural numbers $[1{:}N] \subseteq \mathbb{N}$. Define the predicate $See\_Repeating(j)$ to hold if and only if the repeating item is seen as repeating on the jth trial. Then, we define *first seen as rep. probability* as follows. Eq. (A.1):

$$p_i = p(\, See\_Repeating(i) \mid \forall j \in \mathbb{N}\ (1 \leq j < i) \cdot \neg\, See\_Repeating(j)\,)$$

That is, p_i is the conditional probability that the repeating item is seen as repeating on trial i, given that on all previous trials it has *not* been seen as repeating.

The key property we are interested in is *invariance to number of repetitions*, which holds if and only if, $\forall i, j \in \mathbb{N}\ (1 \leq i, j \leq N) \cdot p_i = p_j$.

Appendix 1.2

Model

We have also implemented a very simple probabilistic model of the two evidence accumulation cases to confirm our intuition re. the invariance of conditional probabilities. Our objective was not to provide a full mechanistic investigation of evidence accumulation in an RSVP context, but rather to simply confirm our intuition that the invariance to number of repetitions property that is central to our argument does follow from a "simple as possible" interpretation of a lack of evidence accumulation compared to accumulation.

As we did for the mathematical definitions, assume a sequence of N trials in which an item repeats, index these trials with the natural numbers $[1\!:\!N] \subseteq \mathbb{N}$. Our models are defined in terms of the updating of two predicates across these trials:

1) $See(i)$ is $true$ if and only if the repeating item is seen on trial i.

2) $See\_Repeating(i)$ is $true$ if and only if the repeating item is seen *as repeating* on trial i.

Note that the target can be seen, without being seen as repeating, which requires it to have been seen multiple times. As a result, we distinguish between *See*ing and *See*ing as *Repeating*.

The *basic model* is defined by the following rules, where we sample ones and zeros according to Bernoulli distributions, i.e. from $Bern(q)$, where q is the probability of getting a one. Eq. (B.1):

$$See(1) = X \sim Bern(p_i)$$

$$See\_Repeating(1) = 0$$

$$\forall i \in \mathbb{N}\ (1 < i \leq N) \cdot$$

$$See(i) = \begin{cases} X \sim Bern(p_i) & , if\ \forall j\ (1 \leq j < i) \cdot \neg See(j) \\ 1 & , \quad otherwise \end{cases}$$

$$See\_Repeating(i) = \begin{cases} 1 & , if\ See(i-1) \vee See\_Repeating(i-1) \\ 0 & , \quad otherwise \end{cases}$$

This basic model can be specialised into the two versions we are interested in by specifying how p_i is calculated.

No Evidence Accumulation is obtained by setting p_i to a constant, say $p\ (0 \leq p \leq 1)$.

Evidence Accumulation is obtained by setting p_i to a monotonically increasing function of i, which does not go out of bounds. The easiest way to do this is to assume simple linear evidence accumulation, where we first define b and C such that, $0 < b < 1$, $0 \leq C < 1$ and $(b.N) + C \leq 1$, in order that accumulation stays within bounds. Then we define the following. Eq. (B.2):

$$p_i = (b.i) + C$$

Note, b is assumed to be sufficiently above zero to exclude seemingly no evidence accumulation patterns being generated.

Other more complex monotonically increasing functions could be used.

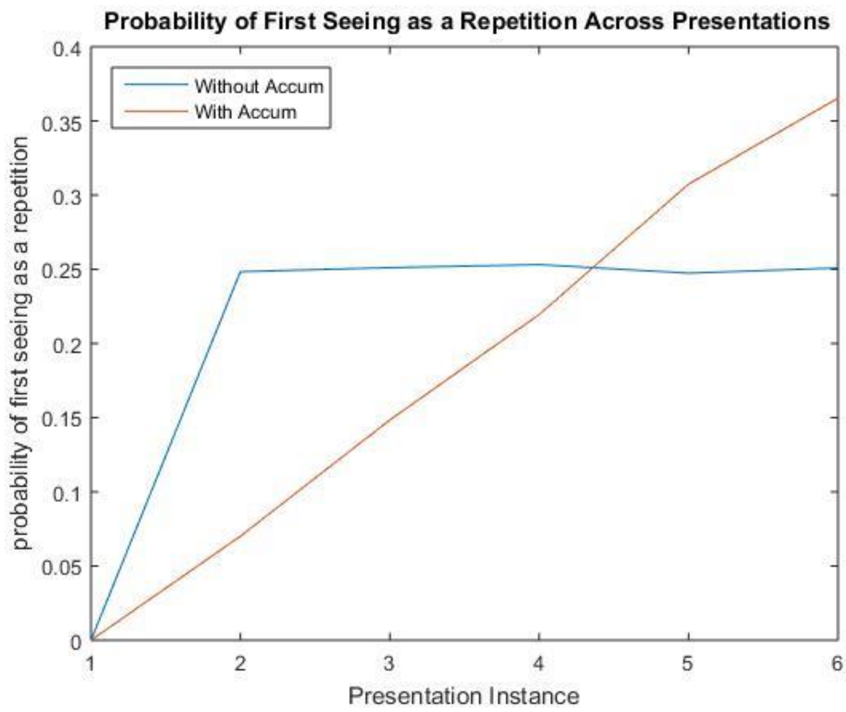

Figure A.1.1: Results of simulations: the fixed probability for the Without Evidence Accumulation case, p in the model description above, was set to 0.25, while $b = 0.075$ and $C = 0.075$ in the With Evidence Accumulation case. The first-seeing as a rep. probability is shown on the y-axis. The With Accum case clearly shows an increasing first-seeing as a rep. probability across repetitions, while the Without Accum case shows no increase in the first-seeing as a rep. probability from the second presentation onwards. Note, a repetition cannot be seen at first presentation, i.e. position one.

Modelling Results

Simulation results are presented in figure 4. The average across 8000 simulated blocks for each of Without and With Evidence Accumulation are presented. Each block reflects six presentations of the repeating item. For each repetition instance, we calculate the first-seeing as a rep. probability, across all the blocks. The fixed probability for the Without Evidence Accumulation case, p in the model description above, was set to 0.25 in the simulations, while $b = 0.075$ and $C = 0.075$ in the With Evidence Accumulation case.

Appendix 1.3

Table A.1.1 Mean probability of first identification as a rep. and (standard deviations). Experiment 1

	Repetition				
SOA	1	2	3	4	5
33	0.02 (0.07)	0 (0)	0 (0)	0 0)	0 (0)
133	0.13 (0.2)	0.22 (0.29)	0.23 (0.35)	0.16 (0.3)	0.18 (0.36)
233	0.29 (0.22)	0.4 (0.34)	0.42 (0.47)	0.42(0.48)	0.19 (0.38)
333	0.33 (0.28)	0.44 (0.39)	0.55 (0.46)	0.67 (0.44)	0.25 (0.50)
433	0.44 (0.33)	0.46 (0.44)	0.27 (0.39)	0.36 (0.46)	0.25 (0.47)
533	0.44 (0.29)	0.69 (0.39)	0.44 (0.47)	0.33 (0.53)	0.12 (0.25)

Table A.1.2 Mean probability of first identification as a rep. and (standard deviations). Experiment 2

	Repetition				
SOA	1	2	3	4	5
17	0.01 (0.07)	0 (0)	0 (0)	0 (0)	0.01 (0.07)
50	0.07 (0.14)	0.07 (0.17)	0.03 (0.1)	0.03 (0.12)	0.04 (0.11)
84	0.08 (0.15)	0.1 (0.17)	0.1 (0.24)	0.09 (0.18)	0.06 (0.16)
117	0.21 (0.22)	0.16 (0.27)	0.27 (0.38)	0.19 (0.38)	0.14 (0.29)
250	0.58 (0.30)	0.35(0.40)	0.43 (0.47)	0.30 (0.49)	0.43 (0.54)
400	0.57 (0.25)	0.60 (0.47)	0.35 (0.48)	0.57 (0.54)	0.33 (0.59)

Contrary to the case of short SOAs, performance at long SOAs (>200ms) increased between the first and second repetition (this pattern is observed in 5 out of 6 long SOA conditions). However, in this analysis, high performance at any repetition entails a loss of power. The underlined probabilities in the table are those where the available data was less than 40% of the data available at the first repetition. When so little data is available, the conditional probability can be considered unreliable. That is, at long SOAs, by the third repetition most participants had identified all the RIs. This emphasizes the fact that when the presentation is clearly above threshold the task is trivially easy. In contrast, at short SOAs participants found it very difficult to identify the RIs, which reinforces the view that it is very difficult to search for items on the basis of episodic information (repetitions) when items are presented at or below threshold.

Appendix 1.4

Table A.1.3 Experimental stimuli (targets) used in Experiments 1 and 2.

Words	**WF**	**AoA**	**Concreteness**	***Familiarity***
List1				
rubber	15	289	596	547
cherry	6	317	611	514
butter	27	206	618	615
muscle	42	397	573	540
cellar	26	361	572	467
cotton	38	306	608	521
avenue	46	372	539	529
fiddle	2	367	582	465
lawyer	43	481	569	520
stable	30	292	562	519
palace	38	294	579	462
timber	19	403	578	440
school	492	228	573	582
racket	5	386	562	486
kettle	3	274	602	551
rabbit	11	206	635	523
kennel	3	322	611	449
hammer	9	278	605	515

Mean	47.5	321.06	587.5	513.61
List1				
pigeon	3	325	609	499
square	143	250	516	576
pocket	46	228	578	590
cavern	1	433	534	400
copper	13	428	547	491
bridge	98	289	623	561
thread	15	267	607	522
orange	23	203	601	567
rattle	5	261	549	448
driver	49	283	553	593
mother	216	144	579	632
banker	5	392	547	524
beggar	2	364	533	435
barrel	24	319	590	487
weapon	42	375	560	517
button	10	192	613	573
singer	10	314	553	548
branch	33	303	583	529
Mean	41	298.33	570.8 3	527.33

WF = Word frequency; AoA = Age of Acquisition.

Appendix 1.5

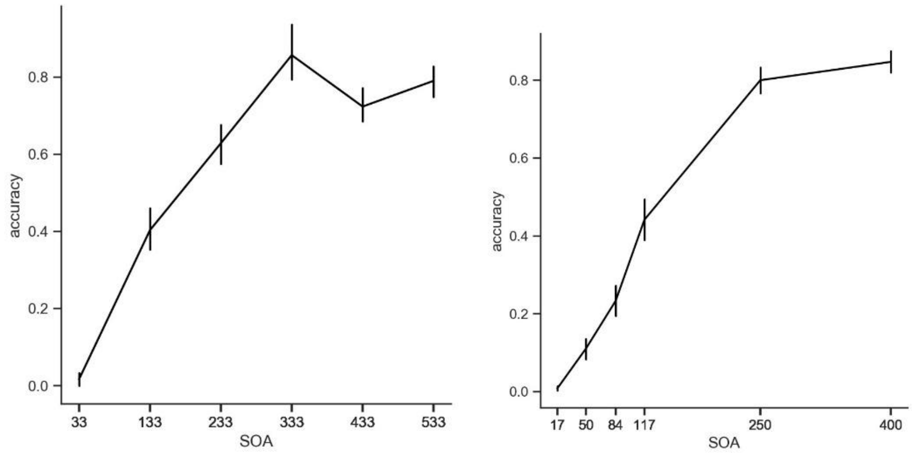

Figure A.1.2. Accuracy (correct identification of the repeated items) for Experiment 1 (left) and Experiment 2 (right).

Appendix 2.1

A potential confound to the inference that recognition changes with SOA more than detection does, is the possibility of a ceiling effect on detection, since this would reduce differences in detection across SOAs, but not differences in recognition (which is far from ceiling). This is unlikely to be the case, since the d' measure is unbounded upwards. Additionally, we present distributions for detection and recognition responses in figure A1. If there were a ceiling effect on detection, their distributions would be narrower than those for recognition. Additionally, a ceiling effect would cause the detection distributions to be skewed, with a longer tail away from ceiling. There is no evidence for either of these characteristics.

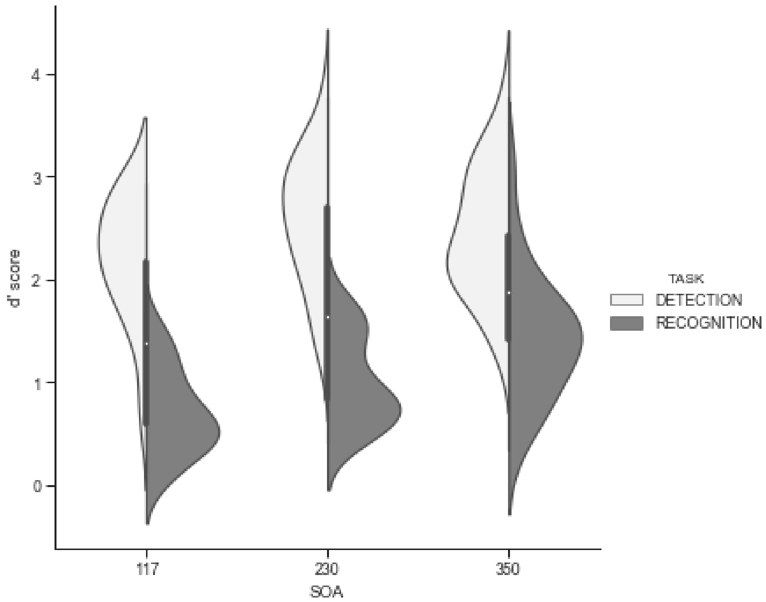

FIGURE A.2

Kernel density estimation of the underlying distribution of the d' values for each SOA condition for Recognition and Detection. The symmetry of the distributions of Detection is not consistent with the skewness that would be expected from a ceiling effect.

Appendix 2.2

REPLICATION

1. Method

1.1. Participants

26 (mean age= 19.4, sd= 0.68) undergraduates students (2 males) participated in the experiment in exchange for course credits. None of them had participated in the main experiment. All were native English speakers and had normal or corrected-to-normal vision. The experiment conformed to the same ethical procedures than the main experiment.

1.2. Materials

Fillers and the targets for detection were the same as in the main experiment. The set of words selected to serve as New Words for the Recognition Task consisted of 144 words.

1.3. Procedure

The screen, stimuli presentation software, and number of trials were the same as in the main experiment. The main difference relative to the main experiment concerns the final item (the mask) in the stream. Four types of masks were used: hashes ("######"), a pseudoword (obtained by substituting two letters of a word), a word or a blank period of time where no stimulus was presented. Each mask was presented in 25% of the streams.

With respect to the questions at the end of the trial there are two types of trials: a) Recognition + detection, and b) Free-recall + detection. (See the main text)

2. Results

Detection task

As in the main experiment, slightly better performance (higher d’ scores) was observed in long SOAs but this time the one-way (SOA: 117, 230, and 350ms) within-subjects ANOVA showed no significant differences ($p>0.1$).

Recognition

We obtained the d’ scores for items presented in each the last three positions (-1, -2, and -3) of the stream before the last masking item in each of the SOA conditions. A 3 (SOA: 117, 230, 350) x 3 (Position: -1, -2, -3) within subjects ANOVA revealed a similar pattern of results as the main experiment. Recognition performance was better in the longer SOAs, $F(2,50)=32.46$, $MSE=11.62$, $p<0.001$, $\eta^2=0.56$ and in the last positions $F(2,50)=10.85$, $MSE=1.71$, $p<0.001$, $\eta^2=0.30$. The interaction was not significant $F(4,100)=0.62$, $MSE=0.07$, $p>0.1$.

Free Recall

The 3 (SOA: 117, 230, 350ms) x 3 (Position: -1, -2, -3) withing subjects ANOVA on the accuracy of the responses revealed a significant effect of SOA $F(2,50)=40.87$, $MSE=0.85$, $p<0.001$, $\eta^2=0.62$, and a significant effect of position, $F(2,50)=18.92$, $MSE=0.36$, $p<0.001$, $\eta^2=0.43$. The interaction did not reach significance, $F(4,100)=1.13$, $MSE=0.01$, $p>0.1$.

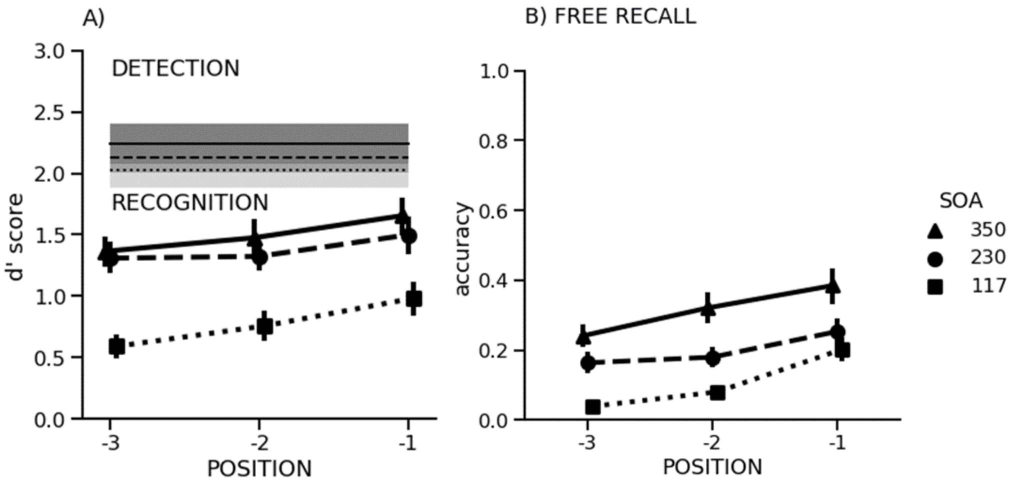

Figure A.2.2. A) Recognition (lower nine points) and detection (upper three lines) performance (left), for three SOAs: solid line 350ms, dashed 230ms, dotted 117ms. Only recognition performance is able to vary with position relative to end of stream: -1, -2, -3. B) Recall performance (right) for the three SOAs and three positions. Error bars and shadowed area indicate standard errors of the mean.

Table 2.1. Mean d' values and (standard deviations) for the Recognition task and mean correct responses and (standard deviations) for the Free-recall task

	Recognition			Free-recall		
SOA	117	230	350	117	230	350
POSITION						
-1	0.98 (0.62)	1.5 (0.68)	1.66 (0.69)	0.2 (0.14)	0.25 (0.14)	0.38 (0.23)
-2	0.76 (0.5)	1.32 (0.52)	1.47 (0.7)	0.08 (0.07)	0.18 (0.11)	0.32 (0.19)
-3	0.59 (0.41)	1.31 (0.56)	1.36 (0.58)	0.04 (0.05)	0.16 (0.13)	0.24 (0.13)

Detection vs. Recognition

A 2 (Taks: Detection, Recognition) x 3 (SOA: 117, 230, 350ms) within-subjects ANOVA revealed a significant effect of Task, $F(1,25)=87.88$, $MSE=32.38$, $p<0.001$, $\eta^2=0.78$. As in the main experiment, performance was better in the Detection task (mean = 2.13, sd = 0.73) than in Recognition (mean = 1.22, sd = 0.61). The effect of SOA was also significant, $F(2,50)=14.47$, $MSE=3.19$, $p<0.001$, $\eta^2=0.37$, where performance was better for longer SOAs. There was a significant interaction between Task and Soa $F(2,50)=5.17$, $MSE=1.12$, $p<0.01$, $\eta^2=0.17$.

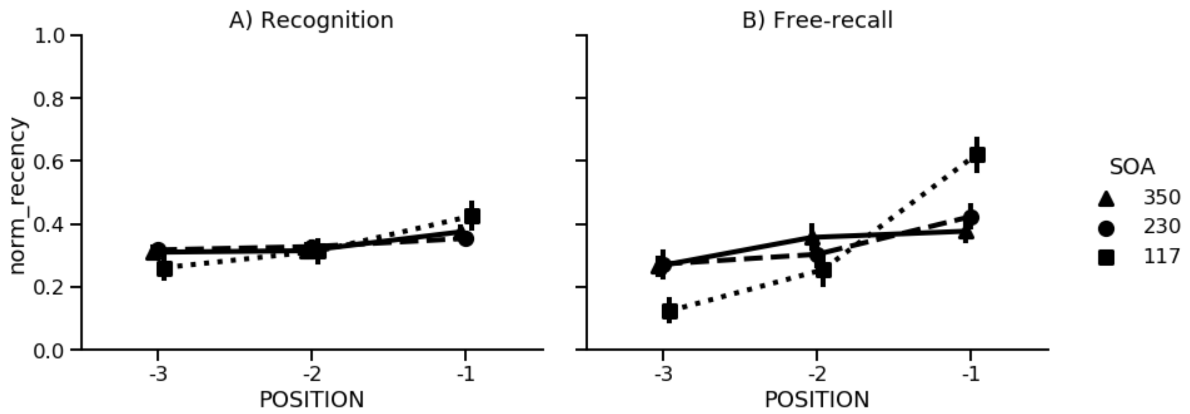

Figure A.2.2. Normalize recency.

Recency

A 2 (Task: free-recall, recognition) x 3 (SOA: 117, 230, 350) x 3 (Position: -1,-2,-3) within-subjects ANOVA[7] showed a significant effect of Position, $F(2,48)=30.47$, $MSE =1.15$, $p<0.001$, $\eta^2=0.34$. Position interacted significantly with task $F(3,72)=6.55$, MSE

---

[7] We excluded the main effects of Task and SOA from the model, since it was singular if either or both were present. This is essentially because the ratios we take for our normalised recency measure, imply that, by construction, there are no differences between levels of these main effects (for example, the sum of all Recognition data points is 3 and the sum of all Recall data points is 3; see figure 4).

=0.17, $p<0.001$, $\eta^2=0.1$ and with SOA $F(6,147)=7.4$, MSE =0.18, $p<0.001$, $\eta^2=0.19$. Finally, the three way interaction approached significance $F(6,147)=2.1$, MSE =0.06, $p=0.057$, $\eta^2=0.08$.

As in the main experiment, the 3 SOA x 3 Position within-subject ANOVAs for each task, revealed a significant interaction in the Free-recall task, $F(6,147)=5.47$, MSE =0.22, $p<0.001$, $\eta^2=0.18$. The same interaction did not reach significance in the Recognition task ($p>0.1$).

Appendix 3

Appendix for Chapter 6

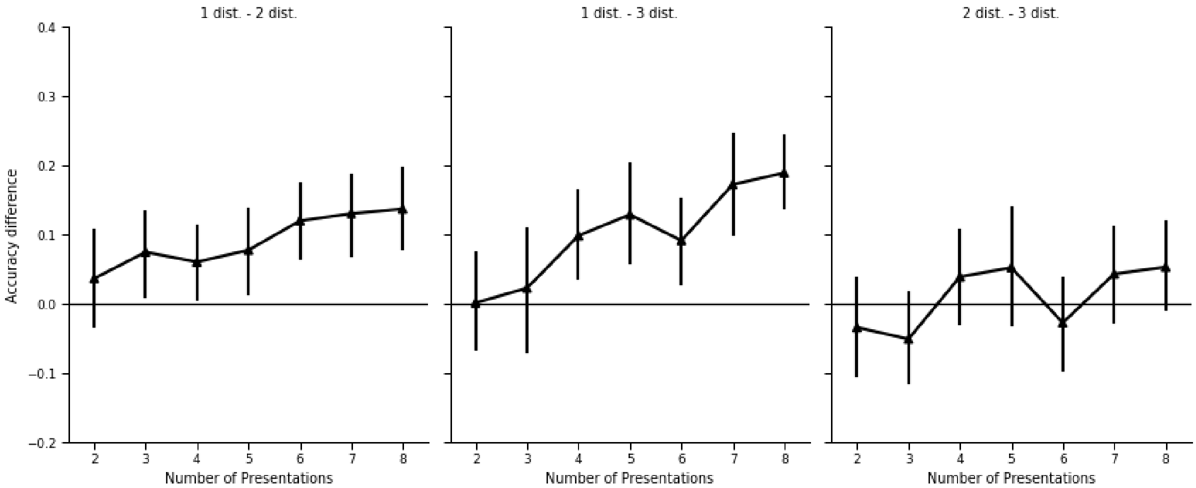

Figure A.3.1. 95% Confidence interval for the differences between conditions in the repetition detection task. One Vs Two-distractors (left), One Vs Three-distractors (centre), and Two Vs Three-distractors (right), across presentations (horizontal axis). Horizontal line at 0 representing no difference between conditions.