Fragile memories for fleeting percepts

Howard Bowman^{1,2} and Alberto Avilés¹

¹ School of Psychology, University of Birmingham, Birmingham, UK.
Edgbaston, Birmingham B15 2TT, United Kingdom
Email: axa1172@student.bham.ac.uk

² School of Computing, University of Kent, Canterbury, UK.
Canterbury, Kent, CT2 7NF, UK
Email: H.Bowman@kent.ac.uk

Correspondence concerning this article should be addressed to Alberto Avilés,
axa1172@student.bham.ac.uk

School of Psychology, University of Birmingham, Birmingham, UK.
Edgbaston, Birmingham B15 2TT, United Kingdom

Abstract

Our perceptual systems are exceptionally good at searching our sensory environments for salient stimuli. A key question is the extent to which this search is performed subliminally. We explore this using Rapid Serial Visual Presentation (RSVP), by comparing detection performance with the memory left for distractors, the stimuli that have to be rejected as non-targets in the process of searching for targets. Our findings are that “immediate” free recall of arbitrary distractors at RSVP rates is very poor, with a severe recency effect. Recognition performance was higher and less subject to recency, but still substantially lower than detection performance. We argue that these findings suggest that the brain subliminally searches for salient stimuli, and are also consistent with a theory we call the tokenized-percept hypothesis, which links conscious perception to the process of episodically marking experiences.

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) (see figure 5 panel A 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. An important 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 very little trace is 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. But, 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*? If rejecting non-targets as not of interest was done 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 reprinted 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.

A key question we explore in this paper is whether the fleeting experiences associated with RSVP, differentially impact recall and recognition. In particular, one might believe that fleeting percepts do not always enable a fully recallable experience of a stimulus, but do, nonetheless, provide a trace in the brain that could support a recognition decision. This resonates with the idea that RSVP streams provide *glimpse percepts*, as suggested by the Glance-Look model of the attentional blink (Su et al, 2011; Barnard & Bowman, 2014), and which might reflect the residual traces left in stage 1 of the Simultaneous Type/ Serial Token (STST) model (Bowman & Wyble, 2007, Bowman et al., 2008).

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 unrepresented words.

2. Method

2.1 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. 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.

2.2 Materials

1700 English words were selected from (Warriner, Kuperman, & Brysbaert, 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.

2.3 Procedure

RSVP streams were presented on a 24'' LCD screen (refresh rate: 60Hz, resolution: 1600 x 1200) using custom PsychToolbox (Brainard, 1997; Pelli, 1997; Kleiner et al, 2007) 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.

3. 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, $F(2,50)=3.50$, $MSE = 0.98$, $p < 0.05$, $\eta^2 = 0.12$. 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, $F(2,50)=17.97$, $MSE = 10.22$, $p < 0.001$, $\eta^2 = 0.43$, and for words presented later in the stream, $F(2,50)=11.77$, $MSE = 2.48$, $p < 0.001$, $\eta^2 = 0.16$. 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, $F(2,50)=14.16$, $MSE = 0.39$, $p < 0.001$, $\eta^2 = 0.31$, and a significant effect of position,

$F(2,50)=27.72$, $MSE =0.81$, $p<0.001$, $\eta^2=0.48$. 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 (0.73)	1.17 (0.53)	1.67 (0.72)	0.23 (0.17)	0.33 (0.2)	0.36 (0.28)
-2	0.65 (0.53)	0.9 (0.65)	1.4 (0.61)	0.05 (0.09)	0.16 (0.17)	0.23 (0.21)
-3	0.56 (0.47)	0.95 (0.63)	1.2 (0.78)	0.05 (0.08)	0.14 (0.16)	0.14 (0.13)

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, $sd=0.57$), $F(1,25)=208.9$, $MSE =71.93$, $p<0.001$, $\eta^2=0.88$. As expected, the main effect of SOA was also significant, $F(1,25)=12.16$, $MSE =3.3$, $p<0.001$, $\eta^2=0.4$. Furthermore, the interaction also reached significance, $F(2,50)=5.5$, $MSE =1.1$, $p<0.05$, $\eta^2=0.18$, 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 A.)

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 ANOVA¹ revealed that the position main effect was significant, $F(2,44)=49.77$, $MSE =2.97$, $p<0.001$, $\eta^2=0.55$. Position interacted significantly with task $F(3,69)=12.45$, $MSE =0.68$, $p<0.001$, $\eta^2=0.3$ and with SOA $F(6,144)=3.11$, $MSE =0.13$, $p<0.05$, $\eta^2=0.14$. Finally, the three way interaction was significant $F(6,135)=2.65$, $MSE =0.1$, $p<0.05$, $\eta^2=0.11$

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, $F(6,135)=4.15$, $MSE =0.21$, $p<0.001$, $\eta^2=0.16$. The same interaction did not reach significance in the Recognition task ($p>0.1$).

¹ 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).

A) Recency operationalised as:

$$\forall j \in \mathbb{Z}(-3 \leq j \leq -1) \cdot \text{norm_rec}_{(j)}(X) = \frac{X(j)}{\sum_{i=-3}^{-1} X(i)}, \text{ with } X \text{ a performance measure}$$

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. $\text{norm_rec}(-1)$ 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)

4. Discussion

We summarise the key findings of the paper, full replications of which are presented in appendix B.

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 4². 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

² We may actually be underestimating recency for both recognition and recall at longer SOAs. This is because differences in memory across SOAs may be underestimated, because there is more time available to forget as SOA increases, due to the time to the memory probe increasing with SOA. This aspect, though, is the same for Recognition and Recall; accordingly it should not confound the comparison of the two.

memory tend to come more into alignment as the involvement of conscious perception increases.

General Discussion

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)³. 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, 2014). Thus, it may be that the fragility

³ Indeed, it is interesting to reflect on the differences between our findings and those with whole-report, where the capacity to encode RSVP items seems good (Nieuwenstein and Potter, 2006). One possible explanation is that whole-report engages a different perceptual regime, one in which the brain is encoding all items, perhaps up to the point at which working memory is full. It may be that the brain enters a new regime once the memory buffer has filled for the first time.

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 (Su et al, 2011; Barnard & Bowman, 2004), 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, e.g. 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.

The Tokenized Percept Hypothesis: The findings here could be argued to bring additional evidence for a further claim, which we have named the tokenized percept hypothesis (Avilés et al, 2020), and is related to the linking of event individuation to awareness by Kanwisher (2001). This hypothesis has emerged from the Simultaneous Type/ Serial Token (STST) model (Bowman & Wyble; 2007) of temporal attention and working memory, where 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.

Avilés et al (2020) 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 Avilés et al (2020) 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.

References:

- Avilés, A., Bowman, H., & Wyble, B. (2020). On the limits of evidence accumulation of the preconscious percept. *Cognition*, 195, 104080.
- 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.
- Barnard, P. J., Scott, S., Taylor, J., May, J., & Knightley, W. (2004). Paying attention to meaning. *Psychological Science*, 15(3), 179-186.
- Bowman, H., & Wyble, B. (2007). The simultaneous type, serial token model of temporal attention and working memory. *Psychological review*, 114(1), 38.
- 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>
- 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).
- Bowman, H., Filetti, M., Alsufyani, A., Janssen, D., & Su, L. (2014). Countering countermeasures: Detecting identity lies by detecting conscious breakthrough. *PloS one*, 9(3).
- Broadbent, D. E., & Broadbent, M. H. (1981). Recency effects in visual memory. *The Quarterly Journal of Experimental Psychology Section A*, 33(1), 1-15.
- Brainard, D. H. (1997) The Psychophysics Toolbox, *Spatial Vision* 10:433-436.
- Cowan, N. (2014). Working memory underpins cognitive development, learning, and education. *Educational psychology review*, 26(2), 197-223.
- Cowan, N. (2016). Working memory capacity: Classic edition. Routledge.
- 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.
- Kanwisher, N. (2001). Neural events and perceptual awareness. *Cognition*, 79(1-2), 89-113.
- Kleiner M, Brainard D, Pelli D, 2007, "What's new in Psychtoolbox-3?" *Perception* 36 ECVF Abstract Supplement.
- McClelland, J. L. (1979). On the time relations of mental processes: an examination of systems of processes in cascade. *Psychological review*, 86(4), 287.
- Nieuwenstein, M. R., & Potter, M. C. (2006). Temporal Limits of Selection and Memory Encoding. *Psychological Science*, 17(6), 471–475. <https://doi.org/10.1111/j.1467-9280.2006.01730.x>
- Pelli, D. G. (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies, *Spatial Vision* 10:437-442.
- 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.
- Potter, M. C. (1976). Short-term conceptual memory for pictures. *Journal of experimental psychology: human learning and memory*, 2(5), 509.

- Potter, M. C., Wyble, B., Haggmann, C. E., & McCourt, E. S. (2014). Detecting meaning in RSVP at 13 ms per picture. *Attention, Perception, & Psychophysics*, 76(2), 270-279.
- 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.
- 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, 348.
- 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>
- 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. <https://doi.org/10.3758/s13428-012-0314-x>
- Wixted, J. T. (2007). Dual-process theory and signal-detection theory of recognition memory. *Psychological review*, 114(1), 152.
- 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.

Appendix A

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.1

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 B

To confirm the reliability of our findings, this appendix contains a full replication of the experiments in the main body of the paper. This replication extends the main-body experiments by including a range of different masks.

METHOD

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 as the main experiment.

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.

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.

There are two types of trials, determined by the end of trial question: a) Recognition + detection, and b) Free-recall + detection. (See the main text)

RESULTS

Trials ending in a blank period of time were not included in the analysis of the Recognition or Free Recall measures, while all trial were used to estimate Detection performance.

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 of 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) within 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 B.1. 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 B.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 (Tasks: 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 B.2. Recency results from the replication experiment for three SOAs and the last three stream positions: A) Normalised recency for recognition, B) Normalised recency for recall.

Recency

A 2 (Task: free-recall, recognition) x 3 (SOA: 117, 230, 350) x 3 (Position: -1,-2,-3) within-subjects ANOVA 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 $F(3,72)=6.55$, $MSE = 0.17$, $p<0.001$, $\eta^2=0.1$ 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 ($p>0.1$).