Consciousness and Cognition 40 (2016) 79–85

Contents lists available at ScienceDirect

Consciousness and Cognition

journal homepage: www.elsevier.com/locate/concog

Consciousness isn't all-or-none: Evidence for partial awareness during the attentional blink

James C. Elliott [a,*], Benjamin Baird [b], Barry Giesbrecht [a]

[a] Department of Psychological and Brain Sciences, University of California, Santa Barbara, United States
[b] Center for Sleep and Consciousness, University of Wisconsin–Madison, United States

ARTICLE INFO

Article history:
Received 1 October 2015
Revised 15 December 2015
Accepted 16 December 2015
Available online 5 January 2016

Keywords:
Attention
Partial awareness
Attentional blink

ABSTRACT

Alternative views of the nature of consciousness posit that awareness of an object is either an all-or-none phenomenon or that awareness can be partial, occurring independently for different levels of representation. The all-or-none hypothesis predicts that when one feature of an object is identified, all other features should be consciously accessible. The partial awareness hypothesis predicts that one feature may reach consciousness while others do not. These competing predictions were tested in two experiments that presented two targets within a central stream of letters. We used the attentional blink evoked by the first target to assess consciousness for two different features of the second target. The results provide evidence that there can be a severe impairment in conscious access to one feature even when another feature is accurately reported. This behavioral evidence supports the partial awareness hypothesis, showing that consciousness of different features of the same object can be dissociated.

© 2015 Elsevier Inc. All rights reserved.

1. Introduction

Pinpointing when and where consciousness occurs in the information processing stream and the nature of this representation are critical steps toward understanding the neurocognitive mechanisms of human awareness. According to proponents of the global workspace theory (Dehaene, Changeux, Naccache, Sackur, & Sergent, 2006) consciousness occurs when an item is available for direct control and manipulation in working memory and this is possible only after extensive information processing. This suggests that consciousness of an object or item is an "all-or-none" phenomenon (Sergent & Dehaene, 2004). However, this view has been challenged by the fact that we seem to be conscious of much more detail than we can report at any given time. In other words, the phenomenology of perceptual experience overflows that which can be accessed (Block, 1995, 2005; Lamme, 2010). The partial awareness hypothesis (Kouider, de Gardelle, Sackur, & Dupoux, 2010) was proposed as a reconciliation between these views and it suggests that consciousness of different levels of representation can occur independently. Therefore, one might be conscious of the color of an object without being able to identify its shape (Breitmeyer, 2014).

In the current study, we used the attentional blink (AB; Raymond, Shapiro, & Arnell, 1992) phenomenon as a tool to compare the all-or-none and partial awareness hypotheses. The AB refers to a decrease in the ability to report a second target object (T2) when it is presented between 200 and 500 ms after a preceding target (T1) in a rapid serial visual presentation

* Corresponding author.
E-mail address: james.elliott@psych.ucsb.edu (J.C. Elliott).

http://dx.doi.org/10.1016/j.concog.2015.12.003
1053-8100/© 2015 Elsevier Inc. All rights reserved.

(RSVP) sequence. The AB is well suited to explore models of consciousness because a stimulus (T2) is presented above sensory threshold but report of the stimulus is often severely impaired. For instance, in a previous investigation, Sergent and Dehaene (2004) examined whether consciousness was an all-or-none phenomenon during the AB. Participants rated the subjective visibility of the T2 stimulus during the AB and it was observed that individuals used only the extreme ends of this scale (i.e., not seen or maximal visibility). This distribution of responses was interpreted as being consistent with an "all-or-none" account of visual consciousness.

Visual masking also can be used to examine the necessary conditions for consciousness. Manipulating the stimulus onset asynchrony (SOA) between a target stimulus and subsequent mask influences the visibility of the target (Bachmann, 2000; Breitmeyer, 1984, 2014). Breitmeyer et al. (2006; see also Breitmeyer, 2014) examined the effect of metacontrast masks on target visibility using two different discriminations. On different trials, participants compared either the luminance or the contour of the target to a sample. It was found that the SOA between the target and the metacontrast mask differently influenced visibility depending on whether the participant was comparing luminance or contour, such that the mask had its peak influence on the contour task at earlier SOAs than in the luminance task (Breitmeyer et al., 2006). Similarly, Hommuk and Bachmann (2009) demonstrated differences in identification and misbinding of the shape and orientation of two stimuli depending on which feature was pre-cued. Both of these experiments suggest that visual masking has distinct influences on different features of the same object. Importantly, though, differences in visual awareness of features doesn't suggest that there is independence in early-stage preconscious processing (Breitmeyer, 2014). To the contrary, there is evidence suggesting that some features, like contour, have a substantial influence on the processing of other stimulus features (Bachmann, 2000; Breitmeyer, 2014). The partial awareness hypothesis extends this work by suggesting that even though parallel unconscious processing of a given feature might depend on the processing of earlier levels or features, the conscious access of a specific feature or level of representation can occur without, or is independent of, conscious access to other levels of representation. Therefore, partial awareness allows for greater richness of perceptual experience at lower levels of the representational hierarchy.

The aim of the present set of experiments is to distinguish between the all-or-none and partial awareness theories of visual consciousness by examining the accuracy of independent discriminations of two features of the same object during the AB. The key assertion of the partial awareness hypothesis is that consciousness occurs for different features and levels of representation independently. Thus, this theory predicts that it should be possible to observe an impairment in accuracy (i.e., an AB) for one T2 feature (e.g., color) even on those trials when another T2 feature (e.g., identity) is accurately reported. In contrast, the all-or-none hypothesis predicts that if there is consciousness of one feature of an object, then one should be able to report other features of the object. The current experiments tested these competing hypotheses by having participants identify two features (color and identity) of T2.

2. Experiment 1

2.1. Methods

2.1.1. Participants

19 undergraduate students from the University of California, Santa Barbara, participated in the experiment for partial fulfillment of a course requirement (11 females, all right handed, mean age = 18.4). Three participants were excluded because of poor performance on the T2 identification task. Based on previous work on the AB (Raymond et al., 1992), a minimum sample size of 16 was set and the final sample size was determined by availability and attendance of participants. The UCSB Human Subjects Committee approved all procedures.

2.1.2. Apparatus and stimuli

Targets and distractors were upper case letters presented on a gray screen in Arial size 32 font ($.51^\circ \times .41^\circ$). T1 was white and T2 was the first non-black (red, green, or blue) letter following T1. All of the distractors prior to the presentation of T2 were black, and those following T2 were either red, green or blue. Stimuli were presented on a 19-in. color CRT monitor positioned 110 cm from the participant. All stimulus presentation, timing, and response acquisition were controlled using the Psychophysics toolbox version 3 (Brainard, 1997).

2.1.3. Design and procedure

Participants initiated each trial by pressing the space bar. Trials began with a 500–1000 ms blank interval during which a fixation cross was on the screen, followed by the RSVP sequence. All items in the RSVP sequence were presented for 96 ms with no interstimulus interval. T1 was always the 10th item. T2 was presented at lag 2, 3, 4, 5, or 9. Both T1 and T2 could be one of 17 letters, excluding B, G, Q, R, V, and W. There were 60 trials per lag, for a total of 300 trials presented in 4 blocks of 75 trials each. After the final RSVP distractor, there was a 500–1000 ms blank interval, followed by the T1 and T2 response screens. Participants were instructed to identify each target by typing the identity using the keyboard and to identify the color of T2 by pressing r, g or b. The order of the response questions was always the same. After the participant responded, the fixation cross reappeared on the screen indicating that the next trial could be initiated when the participant was ready (see Fig. 1).

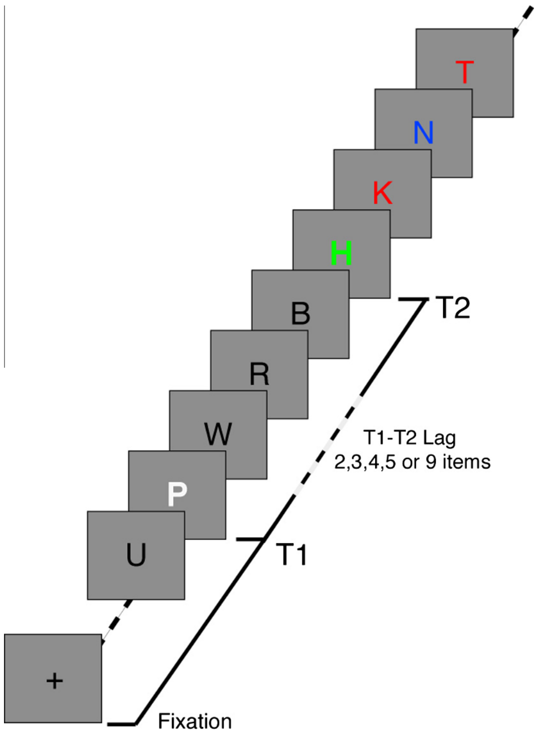

Fig. 1. Schematic representation of Experiment 1 and 2. Each letter in the rapid serial visual presentation (RSVP) was presented for 96 ms. T1 was always the 10th item and was the only white letter in the RSVP. T2 was the first non-black letter following T1.

2.2. Results

The key analysis that contrasts the predictions of the partial awareness and all-or-none hypotheses is whether or not there is an AB for one feature of a stimulus when the other feature is correctly reported. The partial awareness hypotheses, on the other hand, predicts that there should be an AB for one feature even on those trials when the other feature is accurately reported. In order to examine these predictions, two analyses were conducted. The first looked at the accuracy for color on those trials where the identity of the letter was correctly reported. Similarly, the second analysis looked at identity accuracy on those trials where color was accurately reported. There was an AB in the T2 color given accurate identity analysis, $F(4,60) = 4.24$, $p = .014$, $\eta_p^2 = .22$ (see Fig. 2b), such that there was a decrease in accuracy at early lags but not at later lags. Likewise, there was an AB for identity on those trials where color was accurately reported, $F(4,60) = 83.6$, $p < .001$, $\eta_p^2 = .85$ (see Fig. 2a). For comparison, T2 accuracy for color and identity were also plotted independent of whether or not the other feature was accurately reported.

It is possible that during the AB participants were simply guessing on both of the T2 features and that no information regarding either feature was extracted. If true, then the pattern of results in the conditional analysis merely reflects guessing behavior. However, when reporting color (red, green, or blue), performance was significantly better than chance (0.33) at all lags (all p's < .001, Bonferroni corrected critical value = .01). Similarly, when reporting T2 identity, where chance was 1/17 (0.0588), performance was significantly higher than chance (all p's < .01, Bonferroni corrected critical value = .01). Therefore, on average, it is unlikely that participants were merely guessing.

Standard AB analyses were also conducted in order to assess general performance on the task. As expected, T1 accuracy was good overall ($M = 0.801$, SD = 0.14, range: 0.48–0.96). A repeated measures ANOVA with one factor (lag) revealed that lag had a significant effect on T1 performance, such that T1 performance improved as the lag between the targets increased, $F(4,60) = 4.63$, $p < .01$, $\eta_p^2 = .24$. A separate repeated measures ANOVA using T2 color and identity accuracy on T1 correct trials revealed that there was a main effect of feature, $F(1,15) = 363$, $p < .001$, $\eta_p^2 = 0.96$, such that accuracy was higher for reporting color than for identity of the T2. There was also a main effect of lag, $F(4,60) = 121$, $p < .001$, $\eta_p^2 = 0.8$, such that performance improved as the time between the two targets increased. Finally, there was an interaction between lag and feature, $F(4,60) = 11$, $p < .001$, $\eta_p^2 = 0.96$, such that there was a more severe AB for identity compared to color. Follow-up analyses indicated that there was a significant effect of lag for both T2 color accuracy, $F(4,60) = 66.6$, $p < .001$, $\eta_p^2 = .82$, and for T2 identity, $F(4,60) = 115$, $p < .001$, $\eta_p^2 = .89$.

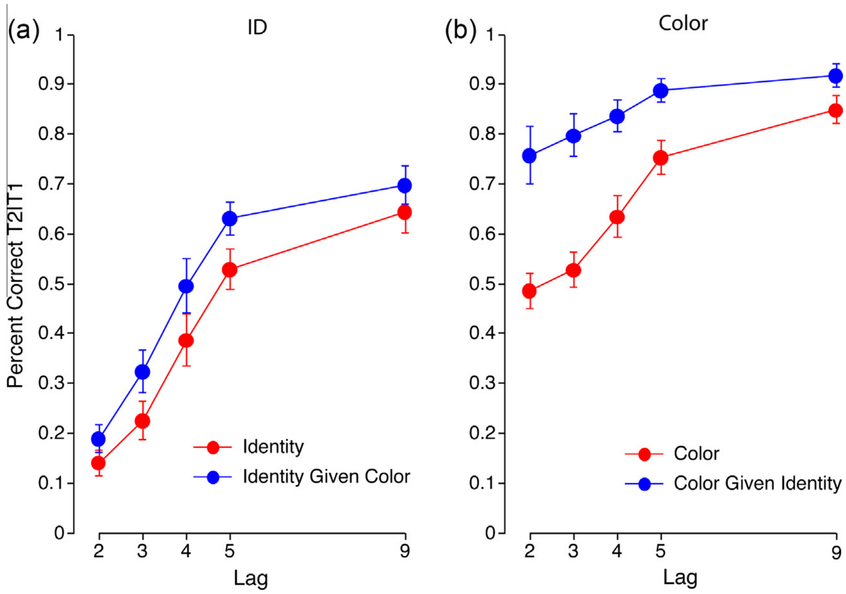

Fig. 2. T2 accuracy in Experiment 1, conditionalized accuracy is plotted in blue for identity (a) and color (b), and for comparison independent accuracy is plotted in red. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

2.3. Discussion

These results support the partial awareness hypothesis. Importantly, since participants performed significantly above chance during the AB in the unconditional analysis, this suggests that sufficient information for an accurate response regarding at least one of the two features was extracted. Therefore, the conditionalized analyses revealed that there was an AB for one feature when the other feature was accurately reported.

Yet, there are asymmetries between the AB observed for color and identity. Specifically, the AB for T2 identity was more severe than the AB for T2 color. This could be due to a number of differences between the selected features. For instance, among other differences, color served as both the target defining characteristic and as a feature to be reported, whereas identity was not a target defining characteristic. However, it is most likely that the largest influence on AB magnitude was the difference in the chance level of performance between the two features. As mentioned, the identity of T2 could be 1 of 17 letters, whereas the color of T2 was always either red, green, or blue. In order to replicate these findings and control for the potential differences due to different levels of chance between the two analyses, experiment 2 constrained the T2 identity to 1 of 3 letters (H, S, or X).

3. Experiment 2

3.1. Methods

3.1.1. Participants

Twenty-six undergraduate students from the University of California, Santa Barbara, participated in the experiment for partial fulfillment of a course requirement (18 females, all right handed, mean age = 18.8). Two participants were excluded because of poor performance in the single target task or poor T1 identification. The UCSB Human Subjects Committee approved all procedures.

3.1.2. Design and procedure

Experiment 2 included two separate tasks: a single target task where participants only reported the color and identity of the T2 item and ignored T1, and a two target task where participants reported T2 color and identity and the identity of T1. In order to equate chance across the two features, the identity of T2 was always either H, S, or X. All other aspects were identical to experiment 1.

3.2. Results

Our primary analysis again evaluated the accuracy of one target feature on those trials when the other feature was accurately reported (conditionalized accuracy). A 2 $\times$ 5 repeated measures ANOVA with task (single target vs. two target) and lag

(2, 3, 4, 5 and 9) as factors was conducted for conditionalized accuracy for each of the features. First, accuracy for color on those trials on which the participants accurately reported identity will be examined. There was a main effect of task ($F(1,23) = 123$, $p < .001$, $\eta_p^2 = 0.842$) such that performance was better in the single target condition ($M = 0.88$, SD = 0.12) compared to the two target condition ($M = 0.73$, SD = 0.135). There was also a main effect of lag ($F(4,92) = 37.6$, $p < .001$, $\eta_p^2 = 0.62$), such that accuracy improved as the time between the two targets increased. Importantly, there was also a significant interaction between task and lag ($F(4,92) = 21.4$, $p < .001$, $\eta_p^2 = 0.482$), such that the time between the two targets influenced accuracy more in the two target condition compared to the single target control, a key indicator of the AB (see Fig. 3d). On those trials when participants accurately reported color, there was a main effect of task on T2 identity accuracy ($F(1,23) = 94.3$, $p < .001$, $\eta_p^2 = 0.80$) and lag ($F(4,92) = 29.6$, $p < .001$, $\eta_p^2 = 0.56$). Finally, there was also a significant interaction between task and lag ($F(4,92) = 12.9$, $p < .001$, $\eta_p^2 = 0.36$), indicating a robust AB for identity even on those trials on which participants accurately reported color (see Fig. 3b). As in experiment 1, accuracy in all conditions was significantly different than chance (.33 for both identity and color, all p's < .001, Bonferroni corrected critical value = .01).

The standard AB analyses were consistent with experiment 1. T1 accuracy was good across all lags ($M = 0.86$, SD = 0.0835, range: 0.63–0.97), but a repeated measures ANOVA with one factor (lag) revealed that T1 performance improved as the time between the two targets increased, $F(4,92) = 3.51$, $p = 0.015$, $\eta_p^2 = 0.132$. A separate repeated measures ANOVA included both lag and T2 features as factors revealed that there was an overall AB, $F(1,23) = 27.3$, $p < .001$, $\eta_p^2 = 0.542$. While there was a main effect of reported feature (color $M = .75$, SD = .11, identity $M = .678$, SD = .143), the interaction between task and lag was not significant (see Fig. 2a and c). Follow-up analyses confirmed that there was an AB for color, $F(4,92) = 64.5$, $p < .001$, $\eta_p^2 = 0.737$, and for T2 identity, $F(4,92) = 55.4$, $p < .001$, $\eta_p^2 = 0.706$.

4. General discussion

The key finding of experiment 1 was that there was still an AB for one feature even when subjects accurately reported the other feature. Analysis of the subset of trials in which a participant accurately reported the color of T2 revealed that there remained a robust AB for the identity of the target. Likewise, on trials in which a participant accurately reported the identity of T2, there remained a robust AB for color. Experiment 2 replicated these results when compared to a classical control condition (single target task) and also demonstrated that the differences in the magnitude of the AB for the different features in experiment 1 was largely due to the differences in chance for each of the features.

The participants' ability to report one feature of a target and not the other demonstrates that awareness of features of objects can occur without awareness of the entire object. The results of these two experiments are therefore consistent with the partial awareness hypothesis (Kouider et al., 2010), which claims that it is possible for a subject to have conscious access of an object at some but not all levels of representation. In contrast, our results provide strong evidence against an all-or-none theory of awareness (e.g. Sergent & Dehaene, 2004), which predicts that when one feature of a stimulus is identified, all of the other features of the stimulus are also identified. However, it is important to note the current experiments do not rule out the possibility that at the level of individual features the transition into awareness might still occur in an all-or-none fashion or occur gradually but extremely fast (Bachmann, 2013). That is, partial awareness may occur on the level of different features of objects, but how a feature transitions into consciousness may or may not be gradual.

The current results contrast with two previously reported experiments. As mentioned, Sergent and Dehaene (2004) found that participants responded using the extremes of a scale when asked about whether or not they were visually aware of T2. This led the authors to conclude that consciousness is all-or-none, as participants never reported intermediate levels of awareness. It is likely that the different conclusions can be explained by the parameters of the two experiments. In the current experiment, participants reported two letters defined by color in a RSVP. In the experiment by Sergent & Dehaene, participants rated the subjective visibility of a number word. In order to successfully rate the subjective visibility of the number word, the participants had to first identify which item in the RSVP was the number word, and then determine the subjective visibility of that word. If participants partially identified a number word, or if different letters or features of the word were conscious to different degrees, then the stimulus would not be identified as a number word. This might cause participants to mistakenly report that they had no visual awareness of the word, when in fact they had partial awareness of the word but not enough to identify it as the T2. This would skew the results of the subjective visibility rating exactly in the direction observed.

More recently, Asplund, Fougnie, Zughni, Martin, and Marois (2014) made similar conclusions using a color discrimination task. Participants searched an RSVP of colored circles for two squares. T1 was always black or white, but T2 was a square that varied in color. Participants responded on a color wheel, allowing Asplund et al. (2014) to determine the precision of the participant's response both inside and outside the AB window. The results showed that there was an overall decrease in report probability during the AB, but that the report precision was constant across lags. The stable precision across lags suggests that the quality of awareness of T2 doesn't fluctuate as a function of lag. This is consistent with the all-or-none hypothesis. A follow-up experiment found similar effects when participants did an analogous task with face stimuli that varied along a pseudo-continuous scale, which suggests that the effect can be observed with complex stimuli. However, given that subjects preferentially select certain features when identifying a face (Peterson & Eckstein, 2013; Sy & Giesbrecht, 2009), it is unclear if these results can be generalized to the entire complex stimulus.

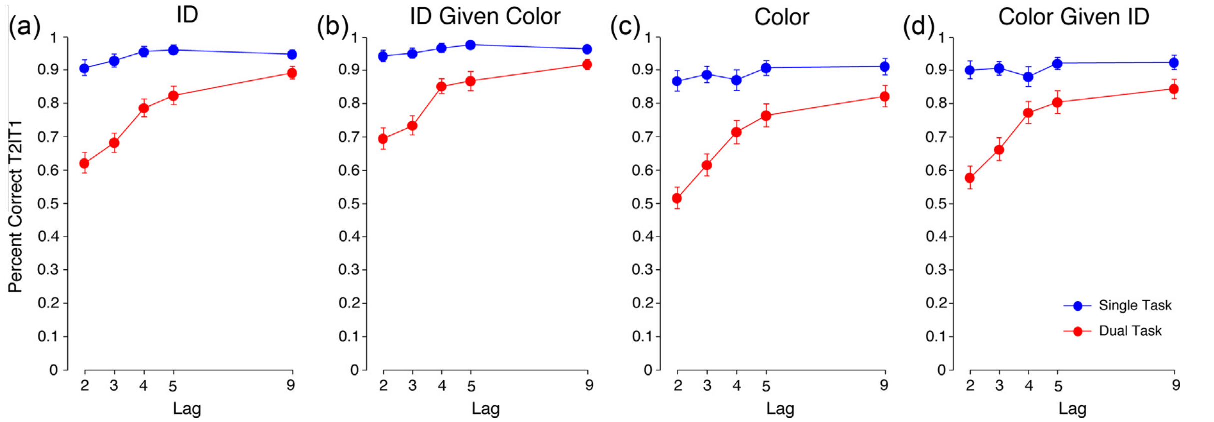

Fig. 3. T2 accuracy for both the single (blue) and dual (red) task conditions. Accuracy for identity (a) and color (c) independent of accuracy on the other feature demonstrates the standard AB. Conditionalized accuracy is plotted for identity given correct report of color (b) and color given correct report of identity (d). (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

As noted above, the current experiments do not rule out the possibility that while awareness of different levels of representation are independent, awareness of any given level could be all-or-none. This possibility was also pointed out in earlier work by Kouider et al. (2010) and, importantly, it is one possible way to reconcile the findings of the current experiment with those of Asplund et al. (2014) and Sergent and Dehaene (2004). Given this explanation, the previous work may be interpreted as showing that there was an all-or-none awareness of certain aspects of an object (e.g., certain letters or features of letters, color or facial features), whereas the current experiments demonstrate that different features of the same object are processed to different levels. The current results converge with the previously discussed work that used visual masking to manipulate awareness (Breitmeyer et al., 2006; Hommuk & Bachmann, 2009) by showing that similar effects are observed when using an attentional manipulation of visual awareness that holds the sensory quality of the stimulus constant.

It is important to note one possible alternative explanation for the results reported here. Specifically, the observed pattern of results could be obtained if T2 identity and/or color were misbound to features of subsequent stimuli. For instance, if on a given trial a subject accurately reported the color of the object and failed to accurately report the identity, it is possible that the identity of the T2 object was bound to the color of one of the subsequently presented letters and that this was the letter that was reported (Chun, 1997). Based on this interpretation, there may have been awareness of both T2 features. However, in experiment 2 the number of possible identities for T2 was limited to H, S, or X and none of the subsequent items in the stream were ever H, S, or X. On those trials when a participant accurately reported the color, it would be impossible for an error in binding to occur because an error in binding would require that the color be bound to the identity of a letter that could not possibly be the second target. Therefore, while the other conditions may represent a combination of both partial awareness and binding errors, this condition unambiguously demonstrates awareness of one feature (color) and a lack of awareness for another (identity).[1]

While the current studies were designed to compare theories of awareness, these data also speak to current theories of the AB. Almost all theories of the AB suggest that all information is processed to a relatively late, post-perceptual stage of processing (Dux & Marois, 2009; Martens & Wyble, 2010). Yet there is evidence that suggests that processing of T2 is flexible during the AB (Elliott & Giesbrecht, 2010; Giesbrecht, Sy, & Elliott, 2007; Giesbrecht, Sy, & Lewis, 2009; Sy, Elliott, & Giesbrecht, 2013). Similarly, given that the results of the current experiment show that features within a single stimulus are processed to different levels, it is problematic to assert that all features between different objects within an RSVP are processed to the same level (whether early or late). We note, however, that it remains possible that both features of the stimulus were processed to a late level and that the suppression of the unreported feature occurred at a post-perceptual stage.

Altogether, our results show that consciousness is not an all-or-none process, but that it is possible to be conscious of one feature of an object but not another. This is clear support for the key assertion of the partial awareness hypothesis, that consciousness of different levels of representations can occur independently, and helps to explain the richness of perceptual experience.

Acknowledgments

This work was supported by a grant from the Yoga Science Foundation and by the Institute for Collaborative Biotechnologies through contract W911NF-09-0001 from the U.S. Army Research Office. The content of the information does not necessarily reflect the position or the policy of the Government and no official endorsement should be inferred.

[1] An alternative approach to rule out the possibility that the results were driven by binding errors would have been to examine whether or not the color and/or identity that was reported was that of the T2 + 1 or subsequent items. However, the identity and color of the items following T2 were not recorded, so that analysis is not possible with the current data.

References

Asplund, C. L., Fougnie, D., Zughni, S., Martin, J. W., & Marois, R. (2014). The attentional blink reveals the probabilistic nature of discrete conscious perception. *Psychological Science, 25*(3), 824–831. http://dx.doi.org/10.1177/0956797613513810.

Bachmann, T. (2000). *Microgenetic approach to the conscious mind* (Vol. 25). John Benjamins Publishing.

Bachmann, T. (2013). On the all-or-none rule of conscious perception. 1–2. doi:http://dx.doi.org/10.3389/fnhum.2013.00387/full.

Block, N. (1995). On a confusion about a function of consciousness. *Behavioural Brain Science, 18*, 227–287.

Block, N. (2005). Two neural correlates of consciousness. *Trends in Cognitive Sciences, 9*(2), 46–52. http://dx.doi.org/10.1016/j.tics.2004.12.006 (review).

Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision, 10*, 433–436. PMID: 9176952.

Breitmeyer, B. G. (1984). *Visual masking: An integrative approach*. New York: Oxford University Press.

Breitmeyer, B. G. (2014). *The Visual (un) conscious and Its (dis) contents: A microtemporal approach*. Oxford University Press.

Breitmeyer, B. G., Kafalıgönül, H., Öğmen, H., Mardon, L., Todd, S., & Ziegler, R. (2006). Meta- and paracontrast reveal differences between contour- and brightness-processing mechanisms. *Vision Research, 46*(17), 2645–2658. http://dx.doi.org/10.1016/j.visres.2005.10.020.

Chun, M. M. (1997). Temporal binding errors are redistributed by the attentional blink. *Perception & Psychophysics, 59*(8), 1191–1199.

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. http://dx.doi.org/10.1016/j.tics.2006.03.007.

Dux, P. E., & Marois, R. (2009). The attentional blink: A review of data and theory. *Attention, Perception, & Psychophysics, 71*(8), 1683–1700.

Elliott, J. C., & Giesbrecht, B. (2010). Perceptual load modulates the processing of distractors presented at task-irrelevant locations during the attentional blink. *Attention, Perception, & Psychophysics, 72*(8), 2106–2114.

Giesbrecht, B., Sy, J. L., & Elliott, J. C. (2007). Electrophysiological evidence for both perceptual and postperceptual selection during the attentional blink. *Journal of Cognitive Neuroscience, 19*(12), 2005–2018.

Giesbrecht, B., Sy, J. L., & Lewis, M. K. (2009). Personal names do not always survive the attentional blink: Behavioral evidence for a flexible locus of selection. *Vision Research, 49*(10), 1378–1388. http://dx.doi.org/10.1016/J.Visres.2008.02.013.

Hommuk, K., & Bachmann, T. (2009). Temporal limitations in the effective binding of attended target attributes in the mutual masking of visual objects. *Journal of Experimental Psychology: Human Perception and Performance, 35*(3), 648–660. http://dx.doi.org/10.1037/a0013481.

Kouider, S., de Gardelle, V., Sackur, J., & Dupoux, E. (2010). How rich is consciousness? The partial awareness hypothesis. *Trends in Cognitive Sciences, 14*(7), 301–307. http://dx.doi.org/10.1016/j.tics.2010.04.006.

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

Martens, S., & Wyble, B. (2010). The attentional blink: Past, present, and future of a blind spot in perceptual awareness. *Neuroscience & Biobehavioral Reviews, 34*(6), 947–957. http://dx.doi.org/10.1016/j.neubiorev.2009.12.005.

Peterson, M. F., & Eckstein, M. P. (2013). Individual differences in eye movements during face identification reflect observer-specific optimal points of fixation. *Psychological Science, 24*(7), 1216–1225. http://dx.doi.org/10.1177/0956797612471684.

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, 21*, 653–662.

Sergent, C., & Dehaene, S. (2004). Is consciousness a gradual phenomenon? Evidence for an all-or-none bifurcation during the attentional blink. *Psychological Science: A Journal of the American Psychological Society/APS, 15*(11), 720–728. http://dx.doi.org/10.1111/j.0956-7976.2004.00748.x.

Sy, J. L., Elliott, J. C., & Giesbrecht, B. (2013). Post-perceptual processing during the attentional blink is modulated by inter-trial task expectancies. *Frontiers in Human Neuroscience*. http://dx.doi.org/10.3389/fnhum.2013.00627/abstract.

Sy, J. L., & Giesbrecht, B. (2009). Target–target similarity and the attentional blink: Task-relevance matters! *Visual Cognition, 17*(3), 307–317. http://dx.doi.org/10.1080/13506280802349746.