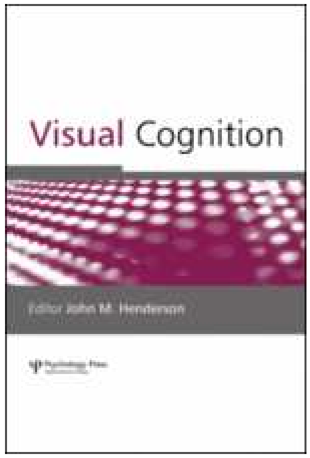

The Cost of Accessing an Object's Feature Stored in Visual Working Memory

Journal:	*Visual Cognition*
Manuscript ID:	Draft
Manuscript Type:	Brief Article
Date Submitted by the Author:	n/a
Complete List of Authors:	Woodman, Geoffrey; Vanderbilt University, Department of Psychology Vecera, Shaun; University of Iowa, Department of Psychology
Keywords:	visual working memory, object-based attention, working memory access

The Cost of Accessing an Object's Feature Stored in Visual Working Memory

Geoffrey F. Woodman[1] and Shaun P. Vecera[2]

Vanderbilt University, Vanderbilt Vision Research Center, Center for Cognitive and Integrative Neuroscience[1]
University of Iowa[2]

Running head: ACCESSING OBJECTS IN VWM

Key words: visual working memory, object-based attention, working memory access

Abstract: 120
Word count: 3,588

Correspondence from the editor/publisher should be addressed to:

Geoffrey F. Woodman
Department of Psychology
Wilson Hall
PMB 407817
2301 Vanderbilt Place
Nashville, TN 37240-1103
615-322-0049 (telephone)
615-343-8449 (fax)
geoffrey.f.woodman@vanderbilt.edu (e-mail)

Abstract

The effects of accessing or retrieving information held in working memory are poorly understood compared to what we know about the nature of information storage in this limited-capacity memory system. Previous studies of object-based attention have often relied upon memory-demanding tasks, and this work could indicate that accessing a piece of information in visual working memory may have deleterious effects upon the other representations being maintained. In the present study, we tested the hypothesis that accessing a feature of an object represented in visual working memory degrades the representations of the other stored objects' features. Our findings support this hypothesis and point to important new questions about the nature of effects resulting from accessing information stored in visual working memory.

Previous studies have suggested that integrated object representations are stored in visual working memory when all of an object's features are task relevant (e.g., Irwin & Andrews, 1996; Lee & Chun, 2001; Luck & Vogel, 1997; Vogel, Woodman, & Luck, 2001; Woodman & Vogel, 2008). In the present study, we examined one of the implications of temporarily storing multiple object representations in memory. Specifically, we investigated how accessing a specific feature of an object representation effects the other stored representations, both of the accessed object and other objects. Our goal was to answer the following question: Is accessing a piece of information from an object independent of previous retrieval operations, or does the process of accessing information stored in working memory depend upon which data were previously accessed?

The general model of working memory posited by Baddeley and colleagues (Baddeley, 1986; Baddeley & Hitch, 1974; Baddeley & Logie, 1999) proposes that the temporary maintenance and manipulation of information is made possible by a multiple component working memory system. The central executive component of this model directs control processes, such as the encoding or manipulation of information in the modality specific slave stores (i.e., the visuo-spatial store and the verbal store). Given such a model, it is possible that performing an operation, like accessing a subset of the information stored in the visual slave store, may interfere with the maintenance of the other information stored in visual working memory. In fact, this is precisely the logic behind the dual-task interference experiments conducted by Baddeley and Hitch and colleagues (Baddeley, 1986; Baddeley & Hitch, 1974) that motivated the working memory model. For example, Quinn and Ralston (1986) had participants perform a spatial movement task (e.g., flapping their arms) to interfere with control processes in visuo-spatial working memory (i.e., encoding and maintenance). In the present study, we examined the possibility that accessing information from one object representation in visual working memory interferes with the maintenance of the unselected representations.

To examine the effects of selectively accessing one datum in visual working memory, we employed an experimental design very similar to that used by Duncan (1984; see also Vecera & Farah, 1994) in a pioneering study of object-based attention. Duncan presented participants with two superimposed objects, a square of variable height with a gap on one side and an oriented, textured line. He required his participants to report either two features of one of these objects (e.g., the height and gap location of the square), or one feature of each object (e.g., the height of the square and orientation of the line). Duncan (1984) found that participants more accurately reported two features from one object than one feature from each object. He interpreted these findings as evidence that attention mechanisms select entire objects and not just locations in space. Awh, Dhaliwal, Christensen, and Matsukura (2001) more recently showed that similar object-based attention effects are found even when the task relies completely upon visual working memory representations (see also Matsukura & Vecera, 2009). Awh et al. (2001) modified Duncan's procedure by informing the participants of the to-be-reported features only after the sample stimulus had been extinguished and masked. These findings provide converging evidence that objects are represented in visual working memory (Irwin & Andrews, 1996; Luck & Vogel, 1997; Vogel et al., 2001; but see, Huang, 2010; Wheeler & Treisman, 2002).

Of primary importance for the current study, Awh et al. (2001) found that the cost for reporting two features from different objects occurred almost entirely when the second response was made. That is, when subjects reported a feature of one object and then switched to reporting a feature of the other object, the accuracy of the second response was significantly lower than the second response in the within-object condition. One interpretation of this finding is that there is a cost of accessing one feature of an object and then accessing a feature of a different object. For example, studies of long-term memory retrieval have suggested that retrieval-induced inhibition may occur when we access information during associative memory paradigms (e.g., Anderson, Bjork, & Bjork, 1994). That is, the act of retrieving an item may interfere with the ability to retrieve other representations from long-term memory. It is possible that the nature of accessing information from visual working memory is one way that this temporary memory system differs from the long-term memory system. Thus, we contrasted the predictions of an account in which accessing an object representation in working memory is detrimental to the other representations being maintained with the predictions of the hypothesis that the visual working memory system exhibits no interference between sequential access operations.

We modified the procedure used by Awh and colleagues (2001) to examine the effects of sequentially accessing different object features from multiple representations stored in visual working memory. As shown in Figure 1, we required participants to sequentially report multiple features from three objects stored in visual working memory. Across trials, we manipulated the order with which participants had to access the stored features. We had participants store information from three objects in memory so that they could not predict on which features they would be tested as the trial unfolded. If we tested their memory for four object features while they maintained information from two, two-feature objects (as in Awh et al., 2001 and Duncan, 1984), then as soon as we tested subjects' memory for the third feature they would have known what the fourth memory probe would be. This would make the data from the sequence of responses difficult to interpret.

If visual working memory uses a read-out mechanism to access a feature representation that is deleterious to the other object representations, then performance should differ between the types of trials in predictable ways. Specifically, on the $\underline{A_1A_2B_1B_2}$ trials (using an $Object_{feature}$ nomenclature, where A and B are just the first and second randomly selected objects) accuracy should be similar for the first two responses because observers are reporting the two features of the same object (similar to Awh et al., 2001; Duncan, 1984). This accessing of one object's features repeatedly should then have compounding negative consequences for the other representations in working memory. Therefore, on $\underline{A_1A_2B_1B_2}$ trials we should observe a large drop in accuracy between the first two and the last two feature reports relative to the $\underline{A_1B_1A_2B_2}$ trials where selective access switches between two objects in memory. In addition, if accessing information has negative consequences for the unselected representations, then the trials in which subjects have to report a feature from the third object (i.e., object C), after first reporting features of two other objects, should be worse than when they are required to switch back to the first two object's we probed (i.e., object A or B). The competing hypothesis we considered is that visual working memory uses an access process that selects and reads out the object features without affecting the other information held in this limited-capacity store. If the latter hypothesis is correct, then accessing the individual pieces of information should be independent of the order in which the objects' features are probed.

Methods

Participants

The participants were 72 undergraduates from Vanderbilt University with normal color vision and normal or corrected-to-normal acuity. They participated for partial fulfillment of a course requirement after informed consent was obtained.

Stimuli

The stimuli were presented using the Psychophysics Toolbox for Matlab (Brainard, 1997) on a homogeneous gray (43.7 cd/m^2) background at a viewing distance of approximately 57 cm. On each trial, participants were shown a red (x = 627, y = 327, chromaticity coordinates of the CIE 1931 color space), blue (x = 142, y = 065), and green (x = 280, y = 589) rectangle. Each rectangle was either thick (0.4° X 1.9° of visual angle) or thin (0.2° X 1.9°) and was tilted at an angle of 45° either to the left or to the right. The three rectangles were centered 6.5° from the white fixation cross (0.4° X 0.4°, 92.6 cd/m^2), each in one of twelve evenly spaced locations on an imaginary circle around the center of the monitor with at least two empty locations between each stimulus. A pattern mask was presented centered over each rectangle's location. The masks were 2.1° X 2.1° and were created individually for each object on each trial by randomly filling the mask's 5 X 5 matrix of cells with red, blue or green. The articulatory suppression load and the memory test response mappings were cued using white letters of a sans serif font (0.4° X 0.5°, 92.6 cd/m^2).

Procedure

Figure 1 depicts an example of the sequence of events that occurred on each trial. Each trial began with a 500-ms presentation of a central fixation cross. Then, the memory array, composed of the red, green, and blue rectangles, was presented for 68 ms. After a 8.5-ms blank interval, the memory items were replaced by pattern masks which remained visible for 200 ms. After the masks were extinguished, a 791-ms retention interval ensued and then the test phase of the trial began (i.e., a 1000-ms stimulus-onset asynchrony between the memory array and the first memory test). A response screen was presented that informed participants which feature of which object they were to report first. The response screen also indicated which keys the participants were required to press to report the feature attribute they remembered for that object on that trial (i.e., 'z' or '/'). Immediately after the first response screen was extinguished, a second response screen was presented that prompted the participant to report a different feature. Each response screen remained visible until the response or 5 seconds elapsed. This procedure continued until the participants had reported four features of two or three object representations held in memory.

The trials only differed in whether they were tested on features of two or three objects and the order of features. The order with which the objects and features were tested was randomized across trials. The mapping of the keys used to report each feature changed randomly across trials but remained the same within a trial.

The participants were given 24 practice trials, followed by six experimental blocks of 24 trials each. After observers initiated each block of trials, they were shown a set of 4 letters or digits that they were required to repeat aloud during each trial of that block to discourage verbal recoding of the visual stimuli. The white alphanumeric characters (either, "a, b, c, d," "w, x, y, z," "1, 2, 3, 4," or "6, 7, 8, 9") where presented for 1500 ms (as in Woodman, Vogel, & Luck, 2001). Participants verbal responses where measured with a microphone connected to a computer synchronized with the stimulus presentation computer and offline analyses verified that they complied with the instructions to perform the articulatory task throughout each trial.

Analysis

Our analyses first focused on the comparisons of the first two responses that had been used in previous object-based attention studies. We entered the accuracy data into an ANOVA with the factors of object probed ($\underline{A_1A_2}$ versus $\underline{A_1B_1}$) and response position (first versus second response). Next, we performed an omnibus ANOVA with the factors of objects probed (A then B, versus, A then B then C), object order ($A_{\underline{1}}A_{\underline{2}}B_{\underline{1}}B_{\underline{2}}/C_{\underline{1}}$, $A_{\underline{1}}B_{\underline{1}}A_{\underline{2}}B_{\underline{2}}/C_{\underline{1}}$, or $A_{\underline{1}}B_{\underline{1}}B_{\underline{2}}A_{\underline{2}}/C_{\underline{1}}$), and response position (1st, 2nd, 3rd, or 4th response). These analyses were collapsed across color because we did not find a significant main effect or any interactions involving the color of the stimuli tested.

Results

Figure 2A shows that accuracy on trials in which two features of one object were probed first (77.6% correct) was higher than when the first two probes were one feature of two different objects (75.9% correct), $F(1,71) = 5.08$, $p < .05$. Response accuracy also significantly decreased with each response (78.7% and 74.8% correct, for the first and second feature reports, respectively), $F(1,71) = 51.31$, $p < .001$. The number of objects probed also interacted with response order due to a greater drop in accuracy for between object feature probes (79.2% and 72.6% correct, for the first and second response, respectively) compared to within object feature probes (78.3% and 76.9% correct, for the first and second response, respectively), $F(1,71) = 16.96$, $p < .001$. These findings replicate the pattern widely observed in the literature (Awh et al., 2001; Duncan, 1984; Kramer, Weber, & Watson, 1997; Vecera & Farah, 1994).

Of primary importance for the hypothesis tested here, we found that the accuracy of feature reports was qualitatively different between $\underline{A_1A_2B_1B_2}$ and $\underline{A_1B_1A_2B_2}$ trials (see Figure 2B, left). The accuracy drop for the third feature reported when sequentially switching between objects A and B was much smaller than when two object-A features were reported then two object-B features. This observation was confirmed with an ANOVA focusing on these two trial types using the factors of object probed (A or B) and response position (1st, 2nd, 3rd, or 4th response). This yielded a significant main effect of response position, $F(3,213) = 52.13$, $p < .001$, and a significant interaction of object probed X response position, $F(3,213) = 3.41$, $p < .05$, due to accuracy of the second response being significantly higher on $\underline{A_1A_2B_1B_2}$ trials than $\underline{A_1B_1A_2B_2}$ trials, $F(1,71) = 13.74$, $p < .001$, and the opposite pattern for the third responses, $F(1,71) = 10.35$, $p < .01$.

In Figure 2B we show accuracy measured on all types of trials in the experiment. These findings provide further support for the hypothesis that accessing a feature of one object reduces the fidelity of the other object representations in working memory. Specifically, we found that the accuracy of feature report was higher when observers could report a feature from an object previously accessed (i.e., A or B, mean = 72.9%) compared to when they had to switch to an object that had not been tested yet on that trial (i.e., Object C, mean = 72.0%). Statistical support for these observations was provided by the ANOVA with the factors of objects probed (A then B versus A then B then C), object order ($A_{\underline{1}}A_{\underline{2}}B_{\underline{1}}B_{\underline{2}}/C_{\underline{1}}$, $A_{\underline{1}}B_{\underline{1}}A_{\underline{2}}B_{\underline{2}}/C_{\underline{1}}$, or $A_{\underline{1}}B_{\underline{1}}B_{\underline{2}}A_{\underline{2}}/C_{\underline{1}}$), and response position (1st, 2nd, 3rd, or 4th response). This ANOVA yielded a significant main effect of object order, $F(2,142) = 3.50$, $p < .05$, due to higher accuracy when participants could sequentially report two features from one object compared to when each feature probed required participants to switch to a different object (mean accuracy of 73.1, 71.3, and 72.9% correct for object sequences of AABB/C, ABAB/C, and ABBA/C, respectively). We also found a significant main effect of response position, $F(3,213) = 103.94$, $p < .001$, because accuracy generally declined across the four responses that were made. Relevant to the patterns shown in Figure 2B, we found significant interactions of object order X response position, $F(6,426) = 2.78$, $p < .05$, and objects probed X object order X response position, $F(6,426) = 4.00$, $p < .001$, largely due to a greater decline in accuracy across the response positions and the object orders when the objects probed included a feature from a third object (i.e., object C). Note that when we probed a feature from the third object (i.e., C), feature report accuracy was significantly lower than when we probed another feature of object A or B in the same response position (*p*s < .05) except the pair-wise comparison of the 4th response position between the ABAB and ABAC orders which went in this same direction but did not achieve significance ($F(1,71) = 1.63$, $p = .20$). Also note that in general the drops in accuracy between the third and fourth responses are smaller than those between the earlier responses in the sequences. A post hoc analysis confirms that the decrease between the first and second response was larger than the difference between the third and fourth across all trial types ($F(1,71) = 17.83$, $p < .001$). This suggests that the deleterious effects of switching between representations may reach an asymptote well above chance levels of performance (i.e., 50% correct).

Finally, we address two possible alternative explanations for these findings that do not require us to consider how information in working memory is being accessed. First, we were sensitive to the possibility that the accuracy effects may have simply been due to participants taking longer to respond with some orderings of the objects probed. This could have been due to either a speed-accuracy tradeoff or because slower reaction times (RTs) caused more time to pass between the presentation of the memory items and when the subsequent features were probed.

Although we found a significant main effect of response position on RT ($p < .05$) due to slower mean RTs with each additional response in the sequence, this factor did not interact with the factors of objects probed or object order in an omnibus ANOVA. These findings rule out the alternative explanations of the accuracy effects based solely on RT differences. Second, it is possible that our effects were due to interference during response selection. Specifically, sequentially reporting values along the same feature dimension might have caused the interference we interpreted as due to switching between objects, because feature repetitions occurred half of the time when switching between objects but none of the time when features of the same object were reported. To address this possibility analytically, we broke down performance for the second through fourth responses to examine the effects of feature switches (e.g., A_{tilt} and B_{tilt} versus A_{tilt} and $B_{thickness}$) in the critical $\underline{A_1A_2B_1B_2}$ and $\underline{A_1B_1A_2B_2}$ trial types. We found that a switch in which feature was reported did not have a significant effect or interact with response position (Fs < 1.0). For example, accuracy of the second report when reporting the same feature from object B that had just been reported from object A was surprisingly similar to accuracy when the reported feature of object B was different from that of object A (A_1B_1 same features: B_1 = 72.7% correct, compared to A_1B_1 different features: B_1 = 72.5% correct). Thus, the effects we observed appear to have been dominated by the nature of accessing the representations in working memory with minimal contributions from response-stage interference.

Discussion

In the present study, we examined how object representations are accessed in visual working memory by requiring participants to report individual features from two-feature objects stored in memory. We found that selectively accessing one object representation in visual working memory led to a decline in accuracy for subsequently reporting the other object representations held in visual working memory. This clearly supports the hypothesis that information is read out of working memory using a mechanism that is detrimental to the unselected object representations held in visual working memory. Our findings address the underemphasized issue of how information is accessed or retrieved from working memory stores (see e.g., Miyake & Shah, 1999).

The results of our experiment have additional implications beyond distinguishing between models of working memory access. The present findings bolster those of Awh et al. (2001) who demonstrated that the original within-object advantage reported by Duncan (1984) can also be observed when participants are not told what features they will need to report until after perceptual processing of the items is complete and the items are in memory. Similar to Awh et al. (2001), our participants were instructed as to what aspects of the stimuli to report long after the perceptual processing had constructed the representations of the objects (i.e., ~800 ms). In addition, the present findings are not due to serial position effects found in scores of previous memory experiments (see Baddeley, 1986), because the memoranda were not presented sequentially. The present findings support proposals that integrated object representations are maintained in visual working memory when all of the features of an object are task relevant (Woodman & Vogel, 2008) and this may explain a number of findings in the object-based attention literature. However, note that other paradigms do lend credence to the idea that perceptual selection may also operate within the framework of object-based representations due to minimal memory demands (Egly, Driver, & Rafal, 1994).

Finally, our study suggests three intriguing accounts of how accessing a piece of information in visual working memory affects the other representations being maintained. First, it is possible that accessing a subset of the information stored in visual working memory is detrimental to the representations that are not accessed via active suppression of the unselected representations. Retrieval mechanisms that lead to the active inhibition of competing representations have been proposed to exist in long-term memory (e.g., Anderson et al., 1994). Second, an existing theoretical framework proposes that accessing representations is an operation based on spatial location (Huang, 2010). This model also proposes that accessing a visual working memory representation is a process performed by visual attention and not a working memory mechanism. Variants of the general procedure used here should be able to help determine whether space- or object-based representations are used when accessing information in visual working memory and whether such access is simply another act carried out by the visual attention mechanism taxed during tasks such as visual search. The third plausible account of our findings is that accessing (or retrieving) an object's feature from visual working memory focuses maintenance mechanisms on that item leaving the unselected items to decay. We believe that future use of the paradigm developed here may be integral in testing the classic hypotheses about of whether forgetting in temporary memory is due to passive decay or active interference between representations (e.g., Baddeley, 2001; J. Brown, 1958; Hartshorne, 2008; Melton, 1963; Peterson & Peterson, 1959; Waugh & Norman, 1965).

Acknowledgements

We would like to thank Sean Polyn, Jeff Zachs, Melissa Beck, and Ed Awh for valuable comments and discussions regarding this study. William Hayward, Liqiang Huang, and two anonymous reviewers provided insightful suggestions that improved that paper. G.F.W is supported by NEI (RO1-EY019882) and NSF (BCS 09-57072) and S.P.V. by NSF (BCS 03-39171).

References

Anderson, M. C., Bjork, R. A., & Bjork, E. L. (1994). Remembering can cause forgetting: Retrieval dynamics in long-term memory. *Journal of Experimental Psychology: Learning, Memory, & Cognition, 20*, 1063-1087.

Awh, E., Dhaliwal, H., Christensen, S., & Matsukura, M. (2001). Evidence for two components of object-based selection. *Psychological Science, 12*, 329-334.

Baddeley, A. D. (1986). *Working Memory*. Oxford: Clarendon.

Baddeley, A. D. (2001). Is working memory still working? *American Psychologist, 56*, 851-864.

Baddeley, A. D., & Hitch, G. J. (1974). Working Memory. In G. H. Bower (Ed.), *The Psychology of Learning and Motivation, Vol. VIII* (pp. 47-90). New York: Academic Press.

Baddeley, A. D., & Logie, R. H. (1999). Working memory: The multiple component model. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control*. Cambridge, UK: Cambridge University Press.

Brainard, D. H. (1997). The Psychophysics Toolbox. *Spatial Vision, 10*, 433-436.

Brown, J. (1958). Some tests of the decay theory of immediate memory. *Quarterly Journal of Experimental Psychology, 10*, 12-21.

Duncan, J. (1984). Selective attention and the organization of visual information. *Journal of Experimental Psychology: General, 113*, 501-517.

Egly, R., Driver, J., & Rafal, R. D. (1994). Shifting visual attention between objects and locations: Evidence from normal and parietal lesion subjects. *Journal of Experimental Psychology: General, 123*, 161-177.

Hartshorne, J. K. (2008). Visual working memory capacity and proactive interference. *PLOS One, 3*(7), e2716.

Huang, L. (2010). What is the unit of visual attention? Object for seleciton, but boolean map for access. *Journal of Experimental Psychology: General, 139*, 162-179.

Irwin, D. E., & Andrews, R. V. (1996). Integration and accumulation of information across saccadic eye movements. In T. Inui & J. L. McClelland (Eds.), *Attention and Performance XVI* (pp. 125-155). Cambridge, MA: MIT Press.

Kramer, A., Weber, T. A., & Watson, S. E. (1997). Object-based attentional selection--Grouped-arrays or spatially invariant representations?: Comment on Vecera and Farah (1994). *Journal of Experimental Psychology: General, 126*, 3-13.

Lee, D., & Chun, M. M. (2001). What are the units of visual short-term memory, objects or spatial locations? *Perception & Psychophysics, 63*, 253-257.

Loftus, G. R., & Masson, M. E. J. (1994). Using confidence intervals in within-subject designs. *Psychonomic Bulletin & Review, 1*(4), 476-490.

Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature, 390*, 279-281.

Matsukura, M., & Vecera, S. P. (2009). Interference between object-based attention and object-based memory. *Psychonomic Bulletin & Review, 16*, 529-536.

Melton, A. W. (1963). Implication of short-term memory for a general theory of memory. *Journal of Verbal Learning and Verbal Behavior, 2*, 1-21.

Miyake, A., & Shah, P. (1999). Models of working memory: An introduction. In A. Miyake & P. Shah (Eds.), *Models of Working Memory: Mechanisms of Active Maintenance and Executive Control* (pp. 1-27). Cambridge, UK: Cambridge University Press.

Peterson, L. R., & Peterson, M. J. (1959). Short-term retention of individual verbal items. *Journal of Experimental Psychology, 58*, 193-198.

Quinn, J. G., & Ralston, G. E. (1986). Movement and attention in visual working memory. *Quarterly Journal of Experimental Psychology., 38*(4-A), 689-703.

Vecera, S. P., & Farah, M. J. (1994). Does visual attention select objects or locations? *Journal of Experimental Psychology: General, 123*, 146-160.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2001). Storage of features, conjunctions, and objects in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 27*, 92-114.

Waugh, N. C., & Norman, D. A. (1965). Primary Memory. *Psychological Review, 72*(2), 89-104.

Wheeler, M., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General, 131*, 48-64.

Woodman, G. F., & Vogel, E. K. (2008). Selective storage and maintenance of an object's features in visual working memory. *Psychonomic Bulletin & Review, 15*, 223-229.

Woodman, G. F., Vogel, E. K., & Luck, S. J. (2001). Visual search remains efficient when visual working memory is full. *Psychological Science, 12*, 219-224.

Figure Captions

Figure 1. Example of the stimulus sequence shown during an individual trial. the key used to report each feature was counterbalanced within subjects.

Figure 2. The findings from the first two responses (A) and the results across all four responses (B) as a function of which features and objects were probed. The error bars represent the 95% within-subjects confidence intervals (as recommended by Loftus & Masson, 1994).

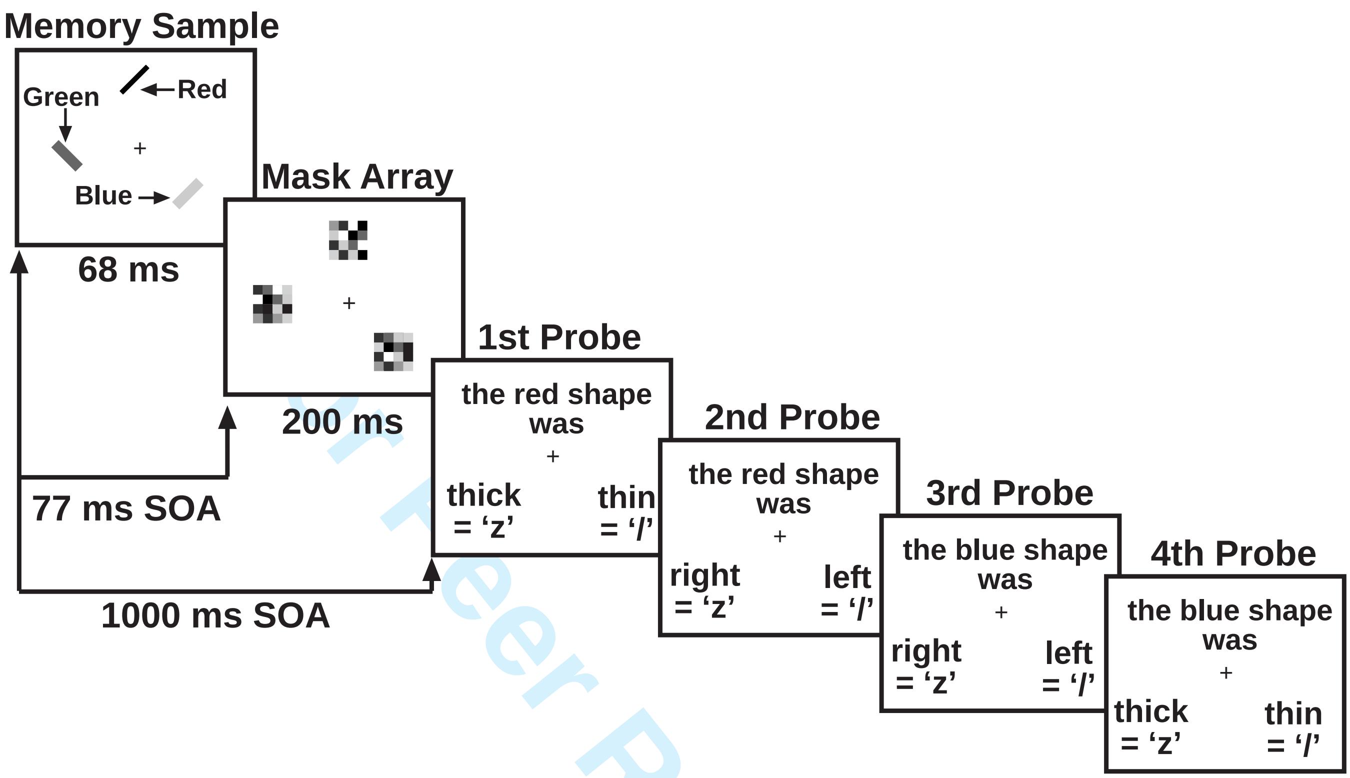

Figure 1. Example of the stimulus sequence shown during an individual trial. the key used to report each feature was counterbalanced within subjects.

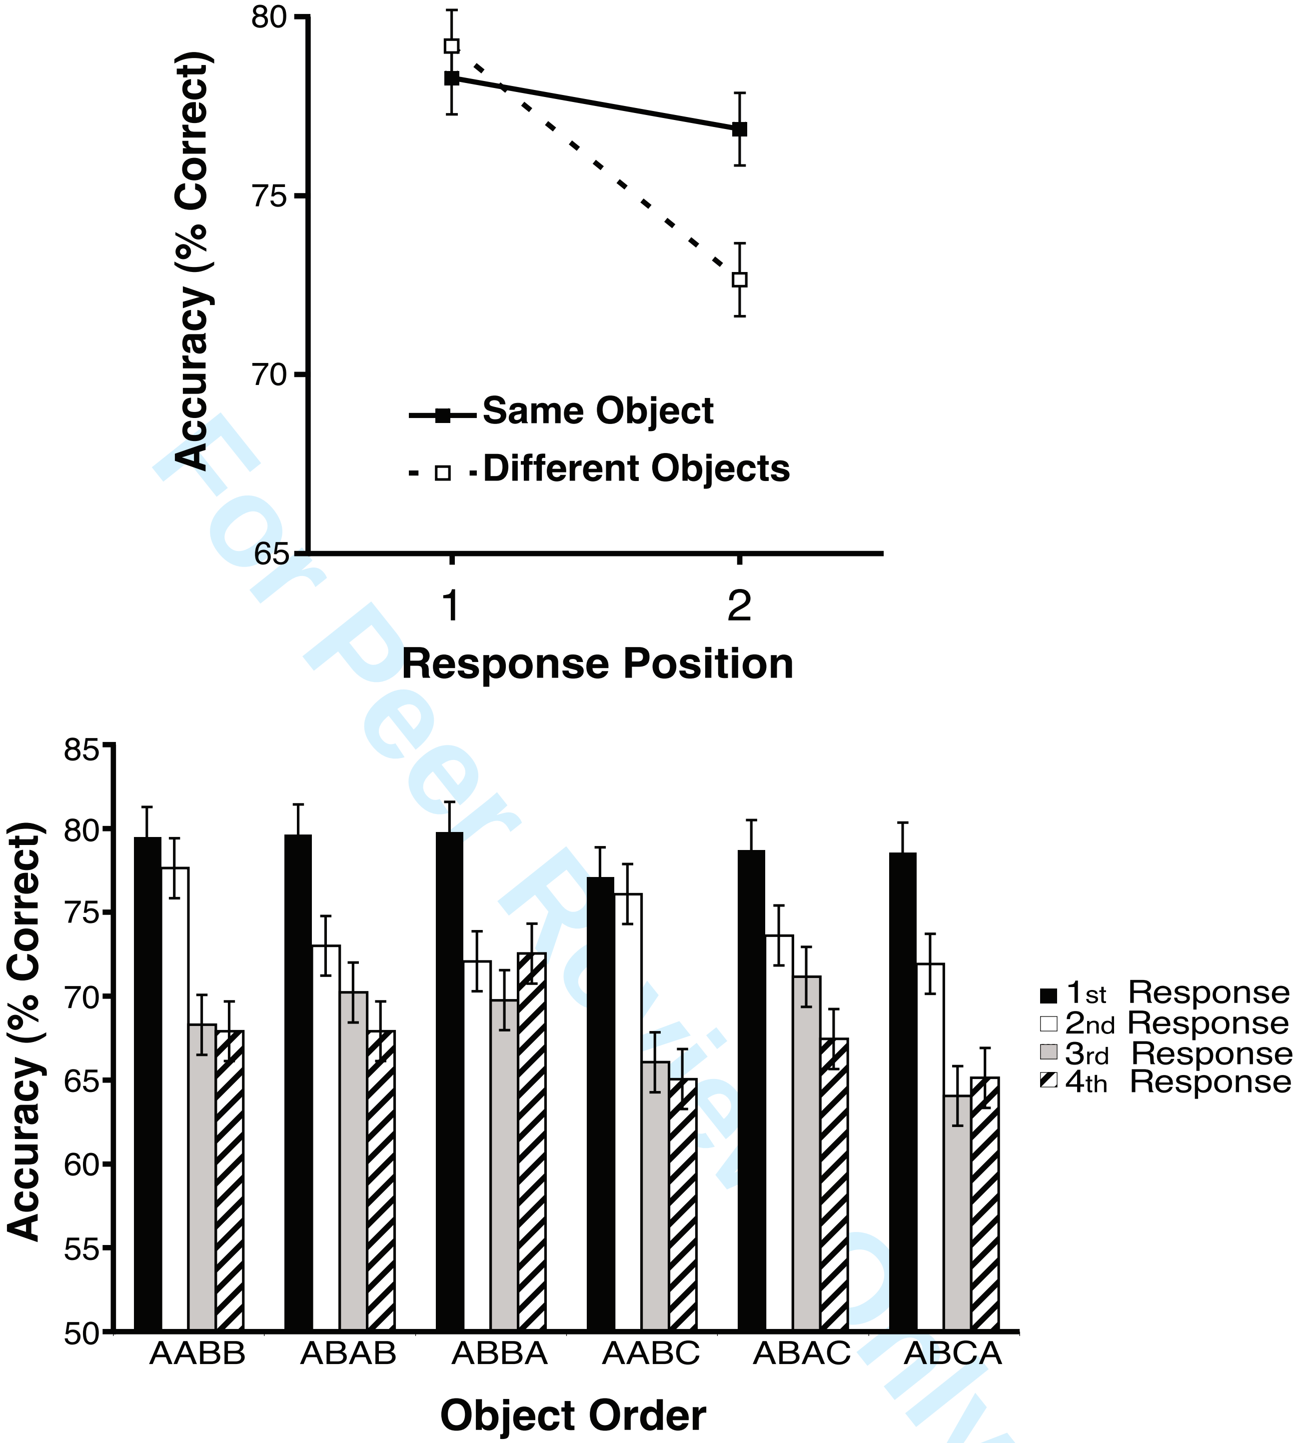

Figure 2. The findings from the first two responses (A) and the results across all four responses (B) as a function of which features and objects were probed. The error bars represent the 95% within-subjects confidence intervals (as recommended by Loftus & Masson, 1994).