Binding in Short-Term Visual Memory: Reassessing Whole Display Interference

Stephen Rhodes¹, Nelson Cowan², Mario A Parra³, & Robert H Logie⁴

¹ Rotman Research Institute

² University of Missouri

³ University of Strathclyde

⁴ University of Edinburgh

This will remain a pre-print. Please cite as:

Rhodes, S., Cowan, N., Parra, M. A., & Logie, R. H. (2019). Binding in Short-Term Visual Memory: Reassessing Whole Display Interference. PsyArXiv. <https://doi.org/10.31234/osf.io/rng4a>

Author Note

Most of this work was completed while the first author was a PhD student at the University of Edinburgh and supported by a studentship from the Centre for Cognitive Ageing and Cognitive Epidemiology (UK cross council Lifelong Health and Well-being Initiative: MR/L501530/1). Part of this work was presented at the 2015 meeting of the Psychonomic Society. Data and code are available at www.github.com/stephenrho/sp-vs-wd.

Correspondence concerning this article should be addressed to Stephen Rhodes, Rotman Research Institute, Baycrest, 3560 Bathurst Street, Toronto, ON M6A 2E1. E-mail: srhodes@research.baycrest.org

Abstract

Here we reassess the finding, originally reported by Wheeler and Treisman (2002), that change detection accuracy for bindings of features is particularly impaired by a whole display probe, in which multiple test items are presented, relative to a single probe. Importantly, the different methods of probing visual short-term memory place different constraints on how information regarding features and their combination can be used to respond. In Experiments 1 and 2 of the present work we use simple processing models of change detection to account for the different task constraints and find no evidence for binding specific whole display interference. Experiment 3 then makes a simple alteration to the single probe task which allows for a more direct comparison of accuracy between the two probing methods and also fails to demonstrate binding specific whole display interference. We argue that the original finding is more reflective of how different change detection tasks allow for different uses of information in memory, rather than reflecting a property of visual short-term memory itself. (168 words)

Keywords: Visual Short-Term Memory; Working Memory; Interference; Cognitive Models

Word count: 9000

Binding in Short-Term Visual Memory: Reassessing Whole Display Interference

The question of how the multiple distinct features of objects are integrated and bound into coherent percepts has been a topic of interest for several decades in cognitive psychology (e.g., Treisman & Gelade, 1980; Wolfe & Cave, 1999). Interest in how bindings of features are retained in visual short-term memory (VSTM)¹ has come about more recently but has already generated a large body of research (e.g., Baddeley, Allen, & Hitch, 2011; Cowan, Blume, & Saults, 2013; Logie, Brockmole, & Jaswal, 2011; Rhodes, Parra, Cowan, & Logie, 2017) (see Schneegans & Bays, 2019 for a review). One of the first papers to directly address the issue of how bound feature representations are maintained in VSTM over brief intervals was by Wheeler and Treisman (2002). This seminal article reported five experiments using different versions of the change detection task to probe memory for individual features (e.g., color, shape), or their precise combination or binding. One crucial finding was that accuracy in detecting changes to the exact binding of features was much lower when VSTM was tested with a *whole display* probe, which contains as many items as initially studied, relative to a *single probe* item (see Figure 1 for examples). In contrast change detection for single features did not differ greatly with the method of probing memory (for similar findings, see Kondo & Saiki, 2012; Yeh, Yang, & Chiu, 2005).

As we describe in more detail below, this finding, and Wheeler and Treisman’s proposed explanation of it, has driven a huge number of studies in the field and a lot of subsequent theorizing. Here we wish to re-examine the basic finding. We note that the two methods of probing VSTM place very different constraints on the use of feature or binding representations in making the same/different or no-change/change decision. These

¹ The terms visual short-term memory and visual working memory are often used interchangeably in this literature. However, there are a number of theories regarding the differences between tasks that require storage only (which some would claim rely on specialized storage mechanisms) and those that require storage plus manipulation of information (in what most would refer to as *working* memory; for reviews see Cowan, 2017; Logie, 2011). We use visual short-term memory here to remain neutral.

differences heavily limit one’s ability to compare the two tasks on accuracy alone. To preview, in our first two experiments we directly compare memory for features and bindings of color-shape stimuli using the two probing methods and use simple item models of change detection to account for the different constraints placed on VSTM by these tasks. In the third experiment we make a simple alteration to the single probe task to facilitate comparisons based on accuracy. Our findings suggest no reason to think that a whole display probe differentially interferes with memory for bindings.

Binding specific whole display interference

In their paper, Wheeler and Treisman (2002) were particularly interested in examining the unit of storage in VSTM. Several years earlier, Luck and Vogel (1997) had found that requiring the storage of multiple features per item did not impair change detection accuracy relative to requiring storage of only a single feature per item (although subsequent experiments have failed to recreate this finding; see Hardman & Cowan, 2015; Oberauer & Eichenberger, 2013). Wheeler and Treisman (2002) noted that Luck and Vogel’s task did not require that the participant retain the *exact* binding of features in memory, as changes involved the introduction of a brand new feature at test. In order to assess such binding they pointed out that it was necessary that the features remain constant between study and test, but the way in which they are *combined* could change (see also, Treisman, 1977). Their experiments 3 and 4 assessed color-location and color-shape binding, respectively, using whole display and single probe change detection tasks. They had separate blocks of trials where a particular feature could change, a block of trials in which either feature could change (the method used by Luck & Vogel, 1997), and a block in which the binding of features could change (see Figure 1A-C for examples). Here we focus on their color-shape experiments, as these are the features we use in the present work (although their findings with color-location were very similar).

In their whole display task, as a change of binding required that two objects swap

features, it was always the case that, even in the blocks assessing memory for individual features, a change occurred to *two* items (e.g., two brand new colors would be introduced in the test array). In the color-shape experiments the locations of items were randomly shuffled between study and test so that location could not influence the decision. In the single probe task the test item was presented in the center of the screen, which was not occupied by any of the studied items, again so that location was uninformative. In the individual feature conditions the single probe either matched a feature (a color or shape) from the studied set or was a brand new feature that was not presented during study. In the binding condition the single probe could match one of the previously presented pairs or it was made up of a color and a shape from two distinct objects from the study set (i.e., a binding change). Wheeler and Treisman (2002) examined these two tasks in separate groups of participants and found that, for whole display change detection, accuracy for binding was lower than that for individual features, whereas with a single probe feature and binding accuracy did not differ significantly. Their conclusion was that “[...] memory for the binding is selectively affected by the perceptual conditions at test, whereas memory for the feature information is not similarly vulnerable” (page 60). They argued that the source of this vulnerability was a reliance of binding on limited attentional resources to maintain the correspondence between parallel feature-specific stores. Processing multiple probe items at test was said to divert attention and result in some degradation of the binding information.

This finding and its interpretation has led to a huge volume of work examining the attention demands of feature binding in VSTM. Many of these studies have been geared towards directly assessing the claim that maintaining feature bindings is more demanding of attention than maintaining individual features. The plurality of these studies have found no evidence that change detection accuracy for bindings is more susceptible to a demanding concurrent task than accuracy for individual features (e.g., Allen, Baddeley, & Hitch, 2006; Allen, Hitch, Mate, & Baddeley, 2012; Johnson, Hollingworth, & Luck, 2008; C. C. Morey &

Bieler, 2013).² Instead, it has been proposed that representations of feature bindings are fragile and are sensitive to overwriting (e.g., Alvarez & Thompson, 2009; Ueno, Allen, Baddeley, Hitch, & Saito, 2011a; Ueno, Mate, Allen, Hitch, & Baddeley, 2011b). The evidence for this latter proposition is growing and, in applying it to the findings of Wheeler and Treisman, we could surmise that multiple test items somehow overwrite the representations of the exact binding of features, whereas a single probe does not cause such overwriting (or, at least, not to the same degree). This issue is also important in thinking about how we assess the capacity of VSTM with the change detection paradigm, as these findings suggest we may get different answers based on the probe procedure used. Unfortunately, comparing accuracy on the single probe and whole display tasks used by Wheeler and Treisman is not as easy as it first appears and, as we describe below, their findings are likely an artifact of the different constraints placed on VSTM by the tasks, rather than reflective of a property of VSTM itself.

Limitations of comparing accuracy between the probing methods

It is clear that change detection accuracy is reflective of a number of factors. For example, in addition to measuring the ability of a participant to do the task, it may also be reflective of their bias towards a particular response option (see Macmillan & Creelman, 2005; Rhodes, Cowan, Parra, & Logie, 2018). Importantly, change detection accuracy is also reflective of the *constraints* placed on the use of information in VSTM by the different probe methods. The constraints in the more standard versions of the single probe and whole display change detection tasks were outlined by Rouder, Morey, Morey, and Cowan (2011), but in the case of the tasks used by Wheeler and Treisman (2002) it is possible that they may explain much, if not all, of the discrepancy between probing methods in accuracy for

² Although, it should be noted that estimates of the *number* of features and bindings in VSTM, estimated using processing models such as the ones discussed in the present work, can result in a different interpretation of these findings, in which a greater proportion of feature bindings are disrupted by a concurrent task relative to features (for more discussion, see Cowan et al., 2013; Cowan, Saults, & Blume, 2014).

bindings.

Consider the whole display task. In this case the constraints for features and bindings are the same. If a change occurred the observer will notice this if one or both of the changed items are in VSTM. Here we are using *items* to refer to whatever is relevant on that particular trial, be it an individual feature or a combination of features. For the single probe task, however, there are different constraints on the use of features and bindings. For features it is necessary to have all of the items in VSTM to be certain of detecting a change. If the probe does not match a feature in memory, and the array size exceeds capacity for that particular feature, the observer cannot be certain that a change occurred (as the probe could be a feature that was not encoded or retained). Only if the probe feature matches a feature in VSTM can the observer be sure that no-change has occurred. For feature bindings, as the single probe item is made up of the features of two objects, the observer can have in memory either one or both of the feature combinations that have been split up to form the probe and detect the change (see Hardman & Cowan, 2016; H. Zhang, Xuan, Fu, & Pylyshyn, 2010). Thus, the bar set by the single probe task for detecting a change for bindings is much lower than that for features. As a consequence, binding accuracy will be biased upwards relative to feature accuracy for the single probe task (see Cowan et al., 2013). A similar bias does not exist for the whole display task. Given this disparity, we can ask whether the findings presented by Wheeler and Treisman (2002) are indicative of a particular susceptibility of feature bindings to multiple probe items or whether the different constraints posed by the different probing methods distort accuracy to merely give this impression.

Thankfully we are able to construct simple processing models of the change detection task, similar to those proposed for other versions of this paradigm (e.g., Cowan et al., 2005; Pashler, 1988; Rouder et al., 2011), that account for the different limitations placed on information use by the different probe types. Note that in the above description of the tasks we are assuming that there are trials on which the observer has no available information

about the set of presented items (features or bindings). This has been a controversial topic over the past decade (e.g., Bays, Catalao, & Husain, 2009; van den Berg, Shin, Chou, George, & Ma, 2012) but nevertheless these items models of change detection have repeatedly demonstrated their utility when stimuli are highly distinct (Donkin, Nosofsky, Gold, & Shiffrin, 2013; Rhodes, Cowan, Hardman, & Logie, 2018; Rouder et al., 2008). Further, recent evidence using a fine grained recall task strongly supports the assumption that observers retain information from only a subset of items and no information about others (Adam, Vogel, & Awh, 2017). These processing models follow the logic outlined in the paragraph above and are outlined fully in the appendix. In our first two experiments we use these processing models to more fairly compare the single probe and whole display change detection tasks. To follow this up, in a third experiment we modify the single probe task to make its constraints on VSTM for features and bindings more similar to the whole display task, which in turn allows us to compare accuracy on these tasks more directly.

Experiment 1

Method

Participants. Twenty-four participants (aged 18–30, 14 females) were recruited from the student community of the University of Edinburgh. Each participant received payment of £5 for the 45 minute testing session. All of the experiments reported here were approved by the University of Edinburgh Psychology Research Ethics Committee.

Stimuli. Items were presented on a grey background on a 20" LCD computer screen. Memory arrays consisted of either 4 or 6 colored shapes each subtending 1.3° of visual angle at an approximate viewing distance of 60 cm (unconstrained). Each object appeared in one of 8 randomly selected locations on a 3×3 grid, measuring $6.7^\circ \times 6.7^\circ$, surrounding the center of the screen (the middle location was not used in the memory array). Objects in the memory array were constructed by selecting 4 or 6 colors and shapes randomly without replacement from master sets of 8 features. The 8 possible colors were brown, pink, orange,

purple, green, blue, red, and yellow, and the 8 possible shapes were arch, hourglass, plus, star, circle, flag, diamond, and chevron. In the single probe condition the test array consisted of a single item presented at the center of the screen, rendering location uninformative. In the whole display condition the test array contained the same number of items as the initial memory array (see Figure 1B for examples). To render location uninformative in this condition, items were shuffled within the original locations randomly between study and test (cf. Wheeler & Treisman, 2002).

Design and Procedure. The general trial procedure used in all experiments is shown in Figure 1A. Participants initiated each trial by pressing the spacebar on the keyboard. The trial began with a fixation cross presented for 500 ms followed by a 250 ms blank screen before the memory array was presented. The memory array remained visible for 500 ms and was replaced by a blank screen for a 1000 ms retention interval.³ The test array then appeared and remained visible until the participant made a response.

The experiment was split into 6 blocks combining the two versions of the change detection task (single probe and whole display) and three memory conditions (color only, shape only, and color-shape binding), as shown in Figure 1B. In the single probe task participants were instructed that half of the time the test object would have been in the memory array and the other half of the time it would have changed. In individual feature

³ In their study, Wheeler and Treisman presented their study array for only 150 ms. Work conducted since has suggested that encoding of items may not be fully complete at this point (Bays, Gorgoraptis, Wee, Marshall, & Husain, 2011; Rhodes, Parra, & Logie, 2016). In addition, Logie et al. (2011) showed that for study-test intervals of less than 1500 ms, randomly changing location as a task-irrelevant feature was disruptive of change-detection for color-shape binding, suggesting perceptual interference with a memory representation of the items in their original location. We chose 1000 ms for this interval to be close to that used by Wheeler and Treisman (900 ms) and thereby to provide a robust test of whether the disadvantage for a whole display relative to single probe was indeed due to visual interference with the representation in memory. As outlined below, despite these temporal differences, we were able to recreate their key accuracy results.

Figure 1. Trial procedure and probe displays for the current Experiments. (A.) Schematic of the general trial procedure used in the present experiments. Participants studied 4 or 6 colored shapes and then, following a delay, responded to a change detection probe. (B.) Examples of test arrays for Experiment 1. (C.) Examples of test arrays for Experiment 2. (D.) Examples of the dual probe test arrays used in Experiment 3. All test display depict a change trial where the correct response is ‘different’. Not drawn to scale.

conditions (color or shape only) they were instructed that a change would involve the introduction of a brand new feature (color or shape) that had not appeared in the original set, and the task-irrelevant feature would not change (for example, in a color block the

shapes would not change between study and test). In the binding condition they were told that a change trial would involve the recombination of a color and shape that were both in the initial set but not paired in the same object.

For the whole display task, participants were instructed that half of the time the objects in the test array would be identical to those in the memory array, albeit shuffled around, and the other half of the trials would contain two changes. In the individual feature conditions of this task participants were informed that a change would involve the introduction of two brand new features to the test objects, with the task-irrelevant features remaining unchanged. In the binding condition a change involved two shapes swapping colors, thereby creating two new combinations of color and shape (Figure 1B).

Half of our participants completed the single probe task first and the other half the whole display task first. The order of memory conditions was counterbalanced with the constraint that each participant completed the different memory conditions in the same order for each version of the task. Participants were given detailed instructions prior to the experiment with visual examples of change and no-change trials for each of the 6 block types. During the experiment participants were given written instructions informing them of the nature of the block (what probe type and memory condition) along with a verbal description of the type of change they were supposed to look out for. Each block consisted of 8 practice trials and 72 experimental trials. Half of the trials in each block were change trials and half contained no-change, evenly distributed across the two set sizes (4 and 6). Participants responded to the test array by pressing keys labelled “*SAME*” or “*DIFF*” which corresponded to the “*z*” and “*m*” keys on the keyboard, respectively. The experiment procedure was implemented in E-prime (Schneider, Eschman, & Zuccolotto, 2012).

Analysis. Change detection accuracy was analyzed with generalized linear mixed models using the R package, *lme4* (Bates, Mächler, Bolker, & Walker, 2015; R Core Team,

2018).⁴ Specifically, the log-odds of a correct response was modeled with the effects of probe type (single probe, whole display), condition (color, shape, binding), and set size along with their interactions. To account for our repeated measures design the models included random participant intercepts and effects of each of the three factors. For the factor of condition we implemented a specific contrast in the model between the individual feature (color or shape only) and binding conditions to test for binding specific effects. The binding condition was coded -1 and the feature conditions were coded 1/2 so that the resulting coefficient contrasts the average of the feature conditions with the binding condition. An orthogonal contrast between color (coded 1) and shape (-1) was also included in the model. The remaining factors were effects coded (probe type: single probe = -1, whole display = 1; set size: 4 items = -1, 6 items = 1). Fixed main effects and interactions were tested via Wald chi-square tests using the `Anova` function from the `car` package (Fox & Weisberg, 2011). To further assess significant main effects or interactions we examine the model coefficients for the specific contrasts we coded.

The accuracy analysis does not take into account whether a particular trial was a change or no-change trial. To estimate the number of features or bindings in VSTM we need to take this into account. As described in the introduction the whole display and single probe tasks place different constraints on the way in which information can be used to perform change detection. In the single probe case different constraints are placed on features and bindings. Thus we used a set of processing models, based on a “slot” conception of VSTM (see Cowan et al., 2005; Pashler, 1988; Rouder et al., 2011 for similar models), that account for these constraints when estimating the number of features or bindings in VSTM. These models are described in detail in the Appendix and are summarized here.

The fundamental assumption is that the observer has k items (which could be features

⁴ We also made use of the `papaja` (Aust & Barth, 2018), `tidyr` (Wickham & Henry, 2018), and `psych` (Revelle, 2018) packages.

or bindings) in memory. Previous work suggests that this value differs for color, shape, and the binding between color and shape (Cowan et al., 2013) so each model we consider allows for different k values for the different conditions. If the relevant information is in memory the observer makes a correct response (we return to what is relevant below), otherwise the observer must guess whether a change occurred. Guessing could be uninformed or it could take into account the number of items in memory (for example, if the single probe is not recognized and the observer had the majority of items in mind, they may be more inclined to guess *change*). We recently found evidence for such informed guessing in more standard single probe and whole display tasks, using color as the to be remembered feature (Rhodes et al., 2018). Nevertheless, in the present work we fit two versions of each model; one assuming uninformed guessing and one assuming that the uninformed guessing rate (u) is updated to account for the number of items in memory (informed guessing). Finally, in previous work with these models it has been necessary to include an “attention” parameter (a) to account for the prospect of lapses, in which the participant is assumed to not attend to some proportion of trials and respond on the basis of an uninformed guess (Rouder et al., 2008).

Returning to what is *relevant* information to make a response in these tasks. In the whole display task used by Wheeler and Treisman, the participant can detect a change if either one or both of the changed items were in memory. Otherwise, if all items in memory match items in the probe, and the number of items presented exceeds capacity, the observer must guess. For the single probe task when bindings are concerned, as a change involves the recombination of two features from two items, a change can be detected if one or both of those bindings was in memory (as the features are sampled without replacement). This is not the case when features are concerned. Here it is necessary to retain all items in memory to be sure of detecting a change (when capacity is exceeded one cannot be sure that the unrecognized probe was not in the presented set). The models described in the appendix attempt to measure the number of features or bindings in VSTM while accounting for these different constraints on information use in these tasks.

To address the issue of primary interest here, we considered three models: In one model k was assumed to vary by stimulus condition (color, shape, binding) but was fixed to be constant across the different probe types and the different set sizes. In the second model the k values were also allowed to differ by probe type to allow for the possibility that different numbers of items may be available from VSTM when probed in different ways (e.g., due to at-test interference). The final model was the most general and allowed for different numbers of items in memory for each set size in addition to probe type. This is to allow for the possibility that more information could be encoded from larger arrays (although see Rhodes et al., 2018; Rouder et al., 2008) or that differences between probing methods may become more pronounced when more items appear in the whole display probe.

We modeled the frequencies of hit and false alarms as binomial random variables with the probability of success determined by the appropriate processing model. Parameter values were estimated for each individual participant by minimizing the negative log likelihood using the `optim` function in R. The search routine for each model was started from 10 random starting values to avoid local minima. The data and code for these analyses are available at <https://github.com/stephenrho/sp-vs-wd>.

Results and Discussion

Accuracy. Figure 2A presents accuracy for the single probe and whole display tasks. There was a significant main effect of condition, $\chi^2(2) = 116.68$, $p < 0.01$, with the model coefficients showing higher overall accuracy for individual features relative to bindings, 0.35 (SE = 0.04), $z = 8.49$, $p < 0.01$ (note coefficients are on log odd scale), and higher accuracy for color relative to shape, 0.22 (SE = 0.03), $z = 7.22$, $p < 0.01$. Set size was also significant, $\chi^2(1) = 97.89$, $p < 0.01$, with lower accuracy with 6 items relative to 4, -0.30 (SE = 0.03), $z = -9.89$, $p < 0.01$. The main effect of probe type was not significant, $\chi^2(1) = 0.01$, $p = 0.94$.

The interaction between probe type and set size was not significant, $\chi^2(1) = 0.75$, $p =$

0.39. There was a significant condition by set size interaction, $\chi^2(2) = 11.06$, $p < 0.01$, such that the difference between features and bindings was smaller at set size 6 relative to 4, -0.10 (SE = 0.03), $z = -3.24$, $p < 0.01$. The contrast between color and shape did not significantly increase with set size, 0.01 (SE = 0.03), $z = 0.48$, $p = 0.63$. Importantly, the interaction between probe type and condition was significant, $\chi^2(2) = 27.95$, $p < 0.01$. Assessing the model coefficients we find that the contrast between features and bindings was indeed larger for the whole display task than the single probe task, 0.17 (SE = 0.03), $z = 5.15$, $p < 0.01$. The contrast between color and shape was not significantly different from zero, 0.05 (SE = 0.03), $z = 1.59$, $p = 0.11$. Finally, the three-way interaction was not significant, $\chi^2(2) = 4.97$, $p = 0.08$.

Thus, despite slight procedural differences, we have recreated the finding of Wheeler and Treisman that is of most interest to us; the discrepancy in change detection accuracy between features and bindings was greater for the whole display task relative to a single probe. Note, however, that one aspect of our findings is difficult to reconcile with the idea that the whole display causes interference in that there was no evidence that the discrepancy between feature and binding accuracy increased with set size. If anything, the discrepancy seems to get smaller with set size, specifically when comparing shape and binding (see Figure 2A).

Number of features and bindings in memory. As noted above, we considered two classes of k model; one in which guessing was uninformed and one in which guessing was informed by the number of items in memory. In line with our previous findings (Rhodes et al., 2018), informed models did a better job of accounting for change detection responses. Data from only one participant of the 24 was fit best by an uninformed model of the six models considered (3 informed, 3 uninformed). Comparing AIC summed across each class of model we find a clear preference for the informed models relative to uninformed ($\Delta\text{AIC} = 36.76$, SE = 5.00). Thus, we proceeded to compare models within this class.

Figure 2. Accuracy and k model predictions (light grey points) for Experiments 1-3 (see main text for details). Error bars are standard error of the mean.

Comparing the least restrictive model to that which fixed k across set size via a likelihood ratio test we find that the restriction imposed does not lead to a significant decrease in model fit for all 24 participants ($G^2(6) < 8.30$, $p > 0.22$). AIC also preferred the more restrictive model that did not allow k to vary with set size ($\Delta\text{AIC} = 10.71$, $\text{SE} = 0.88$).

Next we considered whether we could fix k to be equal across the two probe types (i.e., only one k estimate each for color, shape, binding). This led to a significant decrease in fit for 14 participants and, overall, AIC favored the models with separate k s across probe types by a small margin ($\Delta\text{AIC} = 2.49$, $\text{SE} = 1.16$). Average parameter estimates for this model are presented in Table 1 and the model predictions for accuracy are presented in grey alongside the data in Figure 2A.

Estimates of the number of features and bindings in memory were submitted to a 2 (probe type) by 3 (condition) mixed models ANOVA. There were significant main effects of probe, $F(1, 23) = 27.78, p < 0.01$, and condition, $F(2, 46) = 102.6, p < 0.01$. Estimates were smaller with a whole display ($M = 1.98$, $SD = 0.56$) relative to single probe ($M = 2.86$, $SD = 0.69$). Also estimates were largest for color ($M = 3.34$, $SD = 0.74$), followed by shape ($M = 2.55$, $SD = 0.65$), and then binding ($M = 1.36$, $SD = 0.43$). Finally, there was an interaction between probe and condition, $F(2, 46) = 4.37, p = 0.02$. Follow up t -tests showed that k estimates differed between single probe and whole display for color, $t(23) = -2.80, p < 0.05$, $d = -1.14$ $[-2.00, -0.26]$, shape, $t(23) = -6.40, p < 0.01$, $d = -2.61$ $[-3.70, -1.49]$, and binding, $t(23) = -3.38, p < 0.01$, $d = -1.38$ $[-2.26, -0.47]$. Comparing the size of the k discrepancy between conditions directly, we find that it is significantly greater in the shape condition relative to the binding condition, $t(23) = -3.07, p < 0.01$, $d = -1.25$ $[-2.12, -0.36]$, whereas color and binding do not significantly differ, $t(23) = -0.84, p = 0.41$, $d = -0.34$ $[-1.14, 0.47]$ (the contrast of color and shape was not significant, $t(23) = 1.96, p = 0.06$, $d = 0.80$ $[-0.04, 1.63]$). Thus, it appears that the difference between the two probing methods was, if anything, largest for shape only.

In summary, when we account for the different constraints placed on VSTM by the different probe methods we do not find evidence of binding specific whole display interference. The second experiment aimed to replicate this finding while addressing a potential objection to experiment 1.

Table 1

Mean estimates (with standard error) of the number of features or bindings in visual short-term memory (k), uninformed guessing rates (u), and attention (a) parameters for Experiments 1 to 3. Note that the data from Experiment 3 were better fit by a model that restricted k across probe types, but estimates are provided for comparison. Also note that only one a parameter was estimated per participant per experiment and was applied to both probe types.

	Experiment 1		Experiment 2		Experiment 3	
Parameter	Single Probe	Whole Display	Single Probe	Whole Display	Dual Probe	Whole Display
k_{color}	3.73 (0.20)	2.95 (0.21)	3.91 (0.25)	2.79 (0.21)	2.78 (0.25)	2.83 (0.21)
k_{shape}	3.20 (0.19)	1.91 (0.14)	3.34 (0.22)	2.19 (0.17)	2.49 (0.22)	2.33 (0.17)
k_{binding}	1.65 (0.13)	1.08 (0.11)	1.65 (0.19)	1.40 (0.18)	1.69 (0.19)	1.29 (0.18)
u	0.61 (0.02)	0.37 (0.03)	0.54 (0.02)	0.34 (0.02)	0.55 (0.02)	0.33 (0.02)
a	0.85 (0.03)		0.80 (0.03)		0.86 (0.03)	

Experiment 2

In the individual feature conditions of Experiment 1 the task-irrelevant feature was presented at test. This is somewhat different to the approach of Wheeler and Treisman (2002), as in their single probe color-shape experiment (4B) the task irrelevant feature was held constant at test (colors were presented in squares and shapes were filled in black), whereas in their whole display experiment (4A) the irrelevant features were left unchanged. It is possible that participants used the task-irrelevant feature to restrict memory search, via the use of bound representations. For example, knowing that the circle was previously colored red leads to greater certainty that there is a new feature present when red appears in a triangle, assuming features are sampled without replacement, which is the norm in change detection studies. This information, if used, would be of greater use in the single probe task

where search can be restricted to a single item in VSTM through the use of the task-irrelevant feature, whereas for the whole display it is still the case that the participant has to search for each test item in VSTM until a change is found. This may account for our finding of a larger disparity between probing methods for individual features, particularly shape, in our model-based analysis.

There have been recent investigations into the use of task-irrelevant location information in the change detection task. Both accuracy and latency data suggest that participants do not make use of location information in the single probe task to restrict memory search and, instead, appear to perform an exhaustive search of VSTM for the probed feature (Chen & Cowan, 2013; Cowan et al., 2013; Gilchrist & Cowan, 2014). However, as noted earlier, Logie et al. (2011) showed that for study-test intervals of 1000ms or less, task-irrelevant information may be available in the bound memory representation. Therefore, we assessed whether we could recreate our initial findings when task-irrelevant features were held constant in probe arrays testing VSTM for individual features.

Method

Participants. Twenty-four new participants (aged 18–30, 13 females) were recruited from the student community of the University of Edinburgh. Each participant received payment of £5 for the 45 minute testing session.

Stimuli, Design, and Procedure. The stimuli, design, and procedure used in Experiment 2 were identical to those of Experiment 1 with the only difference being that in the individual feature conditions the task-irrelevant feature was held constant at test. In the color condition test items were presented in squares (a shape outside of the task set) and in the shape condition items were filled in black (a color outside of the task set). Figure 1C presents some example probe arrays.

Results and Discussion

Accuracy. Figure 2B presents accuracy for Experiment 2 and the same general patterns are apparent from Experiment 1. Once again there was a significant main effect of condition, $\chi^2(2) = 53.53$, $p < 0.01$, with higher overall accuracy for features relative to bindings, 0.30 (SE = 0.05), $z = 6.43$, $p < 0.01$, and higher accuracy for color relative to shape, 0.15 (SE = 0.04), $z = 3.58$, $p < 0.01$. The main effect of set size was also significant, $\chi^2(1) = 85.44$, $p < 0.01$, once again with lower accuracy for 6 items relative to 4, -0.25 (SE = 0.03), $z = -9.24$, $p < 0.01$. There was a tendency for higher overall performance with a whole display probe but the main effect of probe type was not significant, $\chi^2(1) = 3.60$, $p = 0.06$.

Turning to interactions, there was a probe type by set size interaction, $\chi^2(1) = 8.24$, $p < 0.01$, such that the drop in performance between 4 and 6 items was greater for whole display, -0.07 (SE = 0.02), $z = -2.87$, $p < 0.01$. The condition by set size interaction was significant, $\chi^2(2) = 10.64$, $p < 0.01$. This was driven by the contrast between features and bindings, which was smaller at set size 6, -0.10 (SE = 0.03), $z = -3.24$, $p < 0.01$, whereas the color vs shape contrast did not differ by set size, 0.01 (SE = 0.03), $z = 0.22$, $p = 0.82$. The probe type by condition interaction was also significant, $\chi^2(2) = 20.17$, $p < 0.01$, and was driven by a larger discrepancy between accuracy for features and bindings with a whole display relative to single probe, 0.14 (SE = 0.03), $z = 4.47$, $p < 0.01$. The contrast between color and shape did not vary significantly by probe type, 0.02 (SE = 0.03), $z = 0.59$, $p = 0.55$. Finally, the three way interaction was not significant, $\chi^2(2) = 5.19$, $p = 0.07$.

Number of features and bindings in memory. Once again the informed class of models outperformed those in which guessing was assumed to be uninformed. The best fitting model was an informed variant for all but one participant and AIC clearly favored informed models ($\Delta\text{AIC} = 33.10$, SE = 6.93).

Restricting k by set size did not significantly reduce fit relative to the unconstrained model for all but one participant ($G^2(6) = 15.35$, $p = 0.02$, p was between 0.05 and 0.1 for 3

participants, and greater than 0.17 for the remaining 20 participants). Overall, AIC clearly favored the restricted model over the model that freely estimated k ($\Delta\text{AIC} = 9.19$, $\text{SE} = 1.54$).

Next we compared the model with different capacities for single probe and whole display tasks to the further restricted model in which k did not differ between probe types. This restriction reduced the quality of fit for only 7 of the 24 participants ($0.05 < p < 0.1$ for 3 participants). Averaged across participants, AIC only very marginally favored the model with separate k s for each probe type ($\Delta\text{AIC} = 0.34$, $\text{SE} = 1.17$). Thus we go on to examine the parameter estimates of this model (see Table 1), but we note that the evidence for different numbers of items available in memory for the different probe types is fairly weak.

An ANOVA on the k estimates revealed a main effect of probe, $F(1, 23) = 22.49$, $p < 0.01$, with larger estimates in the single probe task ($M = 2.97$, $\text{SD} = 0.75$) relative to the whole display task ($M = 2.13$, $\text{SD} = 0.69$). There was also a main effect of condition, $F(2, 46) = 55.48$, $p < 0.01$, with the largest estimates for color ($M = 3.35$, $\text{SD} = 0.84$), followed by shape ($M = 2.77$, $\text{SD} = 0.73$), and then binding ($M = 1.52$, $\text{SD} = 0.71$). The probe type by condition interaction was also significant, $F(2, 46) = 4.39$, $p = 0.02$.

Follow up t -tests showed that k estimates differed significantly between single probe and whole display for color, $t(23) = -3.65$, $p < 0.01$, $d = -1.49$ $[-2.39, -0.57]$, and shape, $t(23) = -4.46$, $p < 0.01$, $d = -1.82$ $[-2.77, -0.84]$, but *not* binding, $t(23) = -1.09$, $p = 0.29$, $d = -0.44$ $[-1.25, 0.37]$. Comparing the size of the k discrepancy between conditions directly we find that it is significantly greater in the shape condition relative to the binding condition, $t(23) = -2.88$, $p < 0.01$, $d = -1.18$ $[-2.04, -0.29]$, and in the color condition relative to binding, $t(23) = -2.21$, $p < 0.05$, $d = -0.90$ $[-1.74, -0.05]$. The color and shape conditions did not differ significantly, $t(23) = 0.08$, $p = 0.94$, $d = 0.03$ $[-0.77, 0.83]$.

In summary, Experiments 1 and 2 replicate the general findings of Wheeler and

Treisman regarding the discrepancy in change detection accuracy between features and bindings; this discrepancy is larger with a whole display relative to a single probe. However, as we argue in the Introduction, comparing accuracy between these probing methods is misleading. When attempting to account for the different constraints placed on visual short-term memory by the different tasks with simple item models of change detection we find that (1) the evidence from model comparison that capacity estimates for bindings and features differ between probe types is quite weak and (2) when we examine these estimates, if anything, the tendency for whole display to produce lower performance is larger for features than bindings (in particular shape in Experiment 1).

One potential issue in comparing these estimates of the number of features or bindings in VSTM is that they impose different maximum values for k for the different tasks. In the whole display task, one of two items is sufficient to detect a change and thus the maximum obtainable value of k is $N - 1$, where N is the set size (this is also the case for single probe binding), whereas for the single probe individual feature task the maximum possible value of k is limited to N . It is possible that these different maximum values have a tendency to bias estimates of k for individual features downwards for the whole display task relative to a single probe (see Rouder et al., 2011 for a similar point). In any case, it would be preferable, and perhaps more convincing, to find a way of comparing probing methods that does not rely on specific models of the task.

Experiment 3

Our aim in this experiment was to devise a way to more directly compare different methods of probing VSTM for features and bindings in terms of change detection accuracy. As the whole display task always requires two changes to occur (given the two necessary changes in the binding condition that constitute a feature swap) it seems reasonable to present two items in the “single” probe task, with the possibility that *neither or both* of the objects have changed. This can still be considered, in essence, a single probe task, as the

observer has to only make a single judgement (*same* or *different*), but they have two items on which to base this judgement. However, for clarity we will refer to it as the “dual probe” task. Figure 1D demonstrates this dual probe task for individual features and bindings in comparison to the whole display task. This dual probe versus whole display comparison is more appropriate given that, in both tasks and all conditions, the aim is to detect whether two changes or no changes have occurred. Indeed, as we outline in the appendix, the processing model for estimating the number of features or bindings in VSTM with this task is very similar to that for the whole display task. Therefore, any inherent difficulties in performing the whole display task, and making several judgments at test, relative to making a single judgement at test should be apparent in the probability of giving a correct response (i.e., accuracy).

Methods

Participants. Twenty-four new participants (aged 18–35, 15 females) were recruited from the student community of the University of Edinburgh. Each participant received payment of £5 for the 45 minute testing session.

Stimuli, Design, and Procedure. The stimuli and general procedure used in this experiment were almost identical to that used in Experiment 2. Probe displays in conditions assessing VSTM for individual features were presented holding the task-irrelevant feature constant. The only difference was that the single probe task was replaced by a dual probe task in which two items were presented at test (Figure 1D). Probe items were presented to the left and right of the center of the screen and did not overlap with any of the initially present memory items. In this task both tested items were either members of the initial memory set or both were different. This was made clear to participants and visual examples of the different changes they were to look out for were given before starting each block of trials.

Results and Discussion

Accuracy. Figure 2C presents accuracy for the dual probe and whole display tasks. The main effects of condition, $\chi^2(2) = 127.22$, $p < 0.01$, and set size, $\chi^2(1) = 113.52$, $p < 0.01$, were significant, whereas the main effect of probe was not, $\chi^2(1) = 0.64$, $p = 0.42$. Accuracy was higher overall for features relative to bindings, 0.47 (SE = 0.04), $z = 11.24$, $p < 0.01$, and was higher for color relative to shape, 0.10 (SE = 0.04), $z = 2.42$, $p < 0.05$. Participants were also less accurate with 6 items relative to 4, -0.32 (SE = 0.03), $z = -10.65$, $p < 0.01$.

The interaction between probe type and set size was not significant, $\chi^2(1) = 0.22$, $p = 0.64$. There was a significant condition by set size interaction, $\chi^2(2) = 46.13$, $p < 0.01$. As in Experiments 1 and 2, this was due to the difference between feature and binding conditions being smaller at set size 6 relative to 4, -0.22 (SE = 0.03), $z = -6.73$, $p < 0.01$. The contrast between color and shape did not significantly increase with set size, 0.02 (SE = 0.03), $z = 0.66$, $p = 0.51$.

Of most interest, the interaction between probe type and condition in this experiment was not significant, $\chi^2(2) = 2.41$, $p = 0.30$. Nevertheless, we examined the contrast of features vs bindings and its interaction with probe type, which was not significantly different from zero, 0.05 (SE = 0.03), $z = 1.45$, $p = 0.15$, and approximately a third of the magnitude of the same coefficient in the first two experiments. Finally, the three-way interaction was also not significant, $\chi^2(2) = 4.97$, $p = 0.08$.

To follow up on the non significant probe type by condition interaction we did an additional cross experiment analysis including all three experiments with a contrast comparing the average of Experiments 1 and 2 to Experiment 3, and another contrast comparing experiments 1 and 2 to each other. Indeed, the probe type by condition interaction was significantly reduced in Experiment 3 relative to Experiments 1 and 2, -0.07

(SE = 0.03), $z = -2.55$, $p < 0.05$.

Thus, the simple change we made to the single probe task, which was to give two items that could either be both the same or both different to items presented in the study array, was sufficient to essentially eliminate differences in accuracy between probe types (compare the left and right panels of Figure 2C to the panels above). Note that, despite this change, the number of items in the whole display probe was either twice or three times that in the dual probe task, which goes against the notion that there is something specific about the whole display probe that interferes with memory for bindings. Together with Experiments 1 and 2, these findings suggest that previous reports of whole display interference had much more to do with the constraints of the task rather than revealing something about VSTM itself.

Number of features and bindings in memory. The primary purpose of this experiment was to facilitate comparison of accuracy between tasks. Nevertheless, it is useful to assess estimates of numbers of features and bindings in VSTM. As in Experiments 1 and 2, versions of the models assuming informed guessing outperformed those assuming uninformed guessing ($\Delta\text{AIC} = 14.81$, SE = 3.60, an informed model fit best for 20/24 participants).

Restricting k by set size did not significantly reduce fit relative to the unconstrained model for all but one participant ($G^2(6) = 13.23$, $p = 0.04$, p was between 0.05 and 0.1 for 2 participants, and greater than 0.11 for the remaining 21 participants). Overall, AIC clearly favored the restricted model over the model that freely estimated k ($\Delta\text{AIC} = 5.70$, SE = 0.74).

Further restricting the k estimates to be the same across probe types reduced the quality of fit for only 4 of the 24 participants ($0.05 < p < 0.1$ for 3 participants). Unlike Experiments 1 and 2, AIC also favored the restricted model ($\Delta\text{AIC} = 1.26$, SE = 0.65) if only slightly. k was largest for color only (M = 2.71, SE = 0.18), followed by shape (M =

2.29, $SE = 0.14$), and then binding ($M = 1.41$, $SE = 0.14$).⁵ The predictions of this winning model are presented alongside the accuracy data in Figure 2C.

Although the better fitting model was the one that restricted estimates of k for features and bindings to be the same across probe types we nevertheless conducted an ANOVA on the estimates from model 2, which estimated different k s for the dual probe and whole array tasks (estimates for this model are presented in Table 1 for comparison to Experiments 1 and 2). There was a main effect of condition, $F(2, 46) = 47.76, p < 0.01$, with the largest estimates for color ($M = 2.80$, $SD = 0.69$), followed by shape ($M = 2.41$, $SD = 0.66$), and then binding ($M = 1.49$, $SD = 0.67$). In line with the results of the model comparison, the main effect of probe type, $F(1, 23) = 1.87, p = 0.19$ (dual probe: $M = 2.32$, $SD = 0.66$; whole display: $M = 2.15$, $SD = 0.58$), and the condition by probe interaction, $F(2, 46) = 0.87, p = 0.43$, were not significant.

General Discussion

Wheeler and Treisman (2002) found that change detection accuracy for the precise binding of features was much lower when probed with a whole display relative to a single probe item. Such an accuracy decrement was not present for change detection for individual features, leading to the interpretation that multiple test items selectively disrupt VSTM for feature bindings. However, as we have outlined, the direct comparison of accuracy between these two tasks is misleading. When we implement a measure of the number of features or bindings in VSTM, which accounts for the different constraints on the use of information posed by the different probing methods, we find that the evidence for an influence of probe type is rather weak. A model that allowed estimates of the number of features or bindings available in VSTM to differ by probe type fit only marginally better compared to a model in which estimates were fixed across probe type. In addition, for Experiments 1 and 2 estimates from this model (see Table 1) suggested that the influence of a whole display relative to

⁵ Additional parameter estimates: $a = 0.89$, $SE = 0.03$; $u_{dp} = 0.56$, $SE = 0.02$; $u_{wd} = 0.34$, $SE = 0.03$.

single probe was *smaller* for bindings than for features. Given concerns about comparing estimates of the number of features or binding in VTSM between conditions with different upper bounds for estimates of the number of features in memory, in Experiment 3 we used a dual probe in which either both or none of the probe items differed between study and test. This task provides a much better match to the demands of the whole display task. This manipulation essentially eliminated differences between probe types, in particular the difference in binding change detection accuracy, suggesting that the constraints of the task play a much more important role in determining accuracy than the number of items in the probe display. Overall these findings are difficult to reconcile with the initial interpretation that the representation of the binding of features in VSTM is particularly susceptible to interfering effects of a whole display probe.

Similarities and differences between testing methods in change detection

Our findings are in line with other previous studies that have compared the more conventional versions of the single probe and whole display change detection tasks and have found more or less identical levels of performance (Jiang, Olson, & Chun, 2000; Rhodes et al., 2018; Vogel, Woodman, & Luck, 2001; Yang, Tseng, & Wu, 2015). It has been suggested that the contextual support afforded by a whole display compensates for any cost of making multiple at-test decisions (Jiang et al., 2000; Makovski, Watson, Koutstaal, & Jiang, 2010). Indeed, in line with this suggestion, there have been findings of whole display superiority when items maintain their locations between study and test (see, e.g., Jiang et al., 2000 Experiment 1). According to this interpretation, our findings may be considered surprising given the fact that contextual information will surely have been distorted when items were shuffled between study and test. Indeed, Jiang et al. (2000) (Experiment 3) compared change detection performance with whole displays in which the locations of items were maintained between study and test to situations where items moved to new locations at test. Changing the locations between study and test resulted in poorer change detection

performance for both color and shape, which Jiang et al. (2000) attributed to the absence of contextual support (although see Logie et al., 2011; Woodman, Vogel, & Luck, 2012 for qualification on this finding). Given that in the present experiments we shuffled item locations between study and test, it is unlikely that whole display performance was boosted by location-based contextual support.

While early conceptions of contextual support from whole display probes suggested that this support was location-based, more recently the importance of ensemble representations in VSTM tasks has been emphasized (e.g., Brady & Alvarez, 2011; H. R. Liesefeld, Liesefeld, & Müller, 2019). Such ensemble representations, which are representations of the higher order properties of the stimuli (such as average hue or luminance of colors in a display), would arguably play a greater role in supporting performance with a whole display probe, where the full set of items can be compared. While the evidence is clear that observers represent higher order properties of visual stimuli (see Brady, Konkle, & Alvarez, 2011 for a review), the influence of such higher order representations in canonical change detection tasks with highly distinct items is not clearly established (for discussion, see Rhodes & Cowan, 2019). Further, while we cannot rule out the idea that whole display performance is boosted by higher order representations, we think that patterns of performance within the present data make this unlikely. Firstly, if representations of overall properties of the array were informing whole display change detection, we would expect a larger effect of increasing set size for this task relative to a single probe. This is because changing two of six items has a smaller effect on the array as a whole compared to changing two out of four items, such that comparing ensemble properties between study and test would provide less support to change detection performance at set size six. This would result in a consistent probe type by set size interaction, which was not observed here. Secondly, representation of ensemble properties of features would be of little use in determining a change of binding between study and test, as the features remain constant. Thus the above interaction of probe type and set size would be expected to be further modulated by whether features or bindings are relevant to the task,

something we also did not observe. In summary, we see no reason to suspect that there is specific interference posed by a whole display probe which is then offset by the extra information provided by this probe type.

That the dual probe and whole display procedures in Experiment 3 did not differ greatly in terms of accuracy makes sense in the light of other findings relating to probing VSTM. The search process engaged in change detection appears to be rather similar irrespective of the number of items to be evaluated at test. Event related potentials suggest that changes are detected rapidly and attention is oriented to a change at the onset of the whole display test array (Hyun, Woodman, Vogel, Hollingworth, & Luck, 2009). Further, for both single and whole display probes, there appears to be an additional capacity limited search through VSTM (Gilchrist & Cowan, 2014; Hyun et al., 2009). Surprisingly, as noted earlier, participants seem to engage in this exhaustive search in the standard single probe task even when location of the item can be used to restrict memory search (see also, Chen & Cowan, 2013; Cowan et al., 2013). Thus, regardless of whether a single item, two items, or many items must be evaluated at test, it seems that observers perform a full search of VSTM in addition to any change signal that may be immediately available (Hyun et al., 2009). It is therefore perhaps unsurprising that we did not find a clear difference in accuracy between the methods of probing in Experiment 3, despite the different numbers of test items to be evaluated.

Relation to other findings on the fragility of bound representations

At first glance our findings appear to conflict with a number of other studies that have suggested, or have been cited in support of the idea, that representations of feature conjunctions are particularly fragile. For example, Logie, Brockmole, and Vandenbroucke (2009) assessed the effect of repeating the same features or conjunctions in memory arrays every three trials (Experiment 1) or on every trial (Experiment 2) on the accuracy of detecting changes to color-shape binding in a whole display probe. Surprisingly, even when

identical conjunctions were presented on every single trial, there was little to no evidence that participants improved at detecting feature swaps across three consecutive test blocks (see also, Shimi & Logie, 2019). This is certainly in line with VSTM being fragile (see also, e.g., Griffin & Nobre, 2003; Phillips & Christie, 1977) but the data do not support the notion that representation of feature bindings are *disproportionately* fragile relative to individual features. Logie et al. (2009) also included trials on which only the color *or* the shape of their six objects was repeated; given that the participants' task was to detect features swapping location this would be expected to improve performance if learning of the features occurred. This was not the case and their Experiment 2 showed that, if anything, repeating the conjunction of features on every trial led to *better* overall performance relative to repeating only one of the features. A third experiment, using recall of conjunctions probed by location, showed that repetition of conjunctions led to learning across trial blocks (see also, Brady, Konkle, & Alvarez, 2009), whereas repetition of features did not to the same degree. The precise mechanisms underlying these patterns of results are unclear (see Shimi & Logie, 2019 for discussion) but clearly neither would be expected if bound representations were *particularly* fragile relative to individual features.

One finding that is potentially more difficult to reconcile with the present work is the *visual suffix* effect (Allen, Castellà, Ueno, Hitch, & Baddeley, 2015; Ueno et al., 2011a, 2011b). Ueno et al. (2011a) used a similar single probe change detection task to that of Wheeler and Treisman, but on some trials presented a to-be-ignored visual suffix (colored shape) 250 ms after the offset of the memory array. This visual suffix appeared to have a general effect on change detection performance but, when the suffix was made up of features that could plausibly come from the task set, Ueno et al. (2011a) observed a small additional deficit for detecting changes to color-shape binding (see also, Allen et al., 2015; Ueno et al., 2011b). Subsequent work has shown that participants are likely to erroneously recall features from the to-be-ignored item (Allen et al., 2015; Ueno et al., 2011b) and has led to the suggestion that bound representations are more susceptible to overwriting from the visual

suffix (Baddeley et al., 2011).

In light of the visual suffix effect, why did we not find evidence for a specific role of the number of items presented at test in determining change detection performance for feature bindings? In addressing this, it is important to emphasize that the visual suffix effect occurs with a *single* item. It may be the case that the presentation of *any* probe, even a single item, is sufficient to disrupt memory for binding information. This would help to explain the generally low number of feature bindings retained in memory relative to individual features, shown in Table 1 (see also Cowan et al., 2013). However, in order for this to also account for the present results we would have to propose that the magnitude of the suffix effect does not scale with the number of to-be-ignored suffixes (if we are assuming that the probe acts like a suffix, a point we return to below). We are unaware of any direct tests, but this seems unlikely given current theoretical accounts of the suffix effect. Such an account is given by Baddeley et al. (2011), who propose that the need for filtering of to-be-ignored stimuli underlies much of the suffix effect with an additional cost to distinguishing relevant and irrelevant bindings when the filter fails (see also, Hu, Hitch, Baddeley, Zhang, & Allen, 2014). The need for filtering would almost certainly scale with the number of to-be-filtered items. Thus, we propose that probe displays in the change detection task, which are, by definition, relevant to the task at hand, do not function like suffixes (as was suggested by Ueno et al., 2011a), which, by virtue of being irrelevant to the task, pose an additional filtering demand.

Note that we are not claiming that test probes have *no* effect on VSTM representations. For example, Tabi, Husain, and Manohar (2019) recently used a fine grained delayed estimation task and found that presenting a probe item at a random starting value (e.g., an arrow rotated 90 degrees) made recall of the studied value (the presented angle) slightly less precise relative to not starting the probe at a random initial value. Rather our contention is that for the coarse representations required for change detection with highly distinct stimuli the distortion posed by probe items is not enough to influence performance

(as supported by the relatively small effect on recall error found by Tabi et al., 2019).

Conclusion

In summary, we have reassessed the suggestion that VSTM for the precise binding of features is particularly susceptible to interfering effects of a whole display probe relative to a single probe item. Once the inherent differences in the two probing methods are controlled, either through the use of process models of change detection or through simple modifications to the task to better match them, we no longer find that the difference between item and binding information in VSTM is different for single-item versus whole display probes

Appendix

Models for estimating k

Whole Display. In the whole display change detection tasks used here and by Wheeler and Treisman (2002) a change trial introduces two new items (features or bindings) into the probe display. The observer will notice this change if one or both of the changed items are in memory. This occurs with probability (cf. H. Zhang et al., 2010):

$$c = 1 - \frac{(N - k)(N - k - 1)}{N(N - 1)},$$

provided $k < N - 1$, otherwise $c = 1$.

However, if the relevant information is not in memory the participant must guess. As outlined by Rouder et al. (2011), it is possible for participants to use knowledge of the number of items they have in memory to update their base expectation of a change. We recently tested the predictions of such an *informed guessing* account for the whole display task with multiple possible changes and for the single probe task and found that it provided better fit than a model in which guessing was uninformed (Rhodes et al., 2018). The base expectation of a change is the uninformed guessing rate, u . With knowledge of the number of items in memory we can work out the likelihood of missing a change when there truly has

been one. This is equal to the prior expectation of a change multiplied by the probability of missing a change, or $(1 - c)u$. To obtain the informed guessing rate we normalize this value by the probability that no change is detected across both change and no change trials, which gives:

$$g_{\text{wd}} = \frac{(1 - c)u}{(1 - c)u + (1 - u)}.$$

A final assumption that is now standard in item models of change detection is that participants lapse with a certain probability and fail to pay attention on a particular trial (with probability $1 - a$). For trials on which a lapse occurred, guessing is assumed to be uninformed. All the above leads to the following expectation for hit rate:

$$h_{\text{wd}} = a(c + (1 - c)g_{\text{wd}}) + (1 - a)u.$$

For trials in which no change occurs, a *change* response can only arise due to guessing (see also Pashler, 1988). Therefore,

$$f_{\text{wd}} = ag_{\text{wd}} + (1 - a)u.$$

For the uninformed version of this model one can replace g_{wd} with u .

Single Probe. A central point of this article is that the single probe tasks used by Wheeler and Treisman (2002) place different constraints on the use of information in memory for features and bindings. Consequently, differences in accuracy partly reflect differences in the number of features and bindings that can be stored but also differences in the constraints posed by the tasks. The models described below aim to account for the constraints placed on discrete item representations. It is important to note that the models below apply to the single probe task in which the probe is presented in a neutral location (i.e. one not occupied by another item) or location is rendered otherwise uninformative. When location is informative, a different model is appropriate (see Cowan et al., 2005; Rouder et al., 2011), although it is unclear that participants actually use the probe location when making change detection responses (Cowan et al., 2013; Gilchrist & Cowan, 2014).

Features.

For individual features the appropriate item model for the single probe task is the reverse, with respect to hit and false alarm rate, of the original Pashler (1988) model for the single change whole display task (see Cowan et al., 2013 for discussion). Thus, for this task we start by considering trials in which no change has occurred. The observer makes a correct rejection when the probed item is present in memory, which occurs with probability $d = \min(k/N, 1)$.

If a match is not detected it is possible for participants to inform their guessing (see Rhodes et al., 2018). The way in which this is assumed to be done follows the same logic as above. The prior expectation of a no change trial, $1 - u$, is multiplied by the probability that the probe is not in memory, $1 - d$, assuming no change had occurred. The probability of guessing *change* in this case is,

$$g_{\text{spf}} = \frac{u}{u + (1 - d)(1 - u)}.$$

Thus for features in this single probe task, the probability of making a false alarm is,

$$f_{\text{spf}} = a(1 - d)g_{\text{spf}} + (1 - a)u.$$

For change trials the probability of a hit is,

$$h_{\text{spf}} = ag_{\text{spf}} + (1 - a)u.$$

Bindings.

When a change occurs in the binding version of the single probe task two features from two different objects are recombined to form the probe. If one or both of the objects that make up the probe are in memory then the change is detected. The probability of this occurring is given by c . On the other hand when no change occurs, the participant only

notices this when the probe binding is in memory, as given by d . When neither a partial match (in the case of a change trial) or a full match (in the case of no change) is detected the observer must guess and, once again, it is possible for this guessing to be informed:

$$g_{\text{spb}} = \frac{(1-c)u}{(1-c)u + (1-d)(1-u)}.$$

The prediction for hit rate from this model is,

$$h_{\text{spb}} = a(c + (1-c)g_{\text{spb}}) + (1-a)u,$$

and for false alarms is,

$$f_{\text{spb}} = a(1-d)g_{\text{spb}} + (1-a)u.$$

Once again these models can be made into uninformed versions by setting g_{spf} or g_{spb} equal to u .

Dual Probe. The primary purpose of the dual probe task used in Experiment 3 is to facilitate direct comparison of accuracy with the whole display task. Nevertheless, constructing a processing model for this task and examining parameter estimates is instructive. In this task if one or both of the probe items match items in memory the participant will respond “same” or “no change”. Otherwise the participant must guess. Given knowledge of the number of items in memory the probability that a change has occurred, given that no match is detected, is,

$$g_{\text{dp}} = \frac{u}{u + (1-c)(1-u)}.$$

Thus, on trials where a change has occurred and the array size exceeds k the hit rate is given by,

$$h_{\text{dp}} = ag_{\text{dp}} + (1-a)u.$$

When no change has occurred the observer notices this when one or both items are in

memory (with probability c). Therefore, the prediction for the probability of a false alarm is,

$$f_{\text{dp}} = a(1 - c)g_{\text{dp}} + (1 - a)u.$$

Notice that the model for the dual probe task is essentially the mirror image of the model for the whole display task. Hit and false alarm rate switch roles as task performance logically must primarily rely on non-match detection for the whole display task, whereas it must rely primarily on match detection for the dual probe task. As accuracy aggregates across hit and false alarm rates it can more fairly be used to compare these two tasks.

In all of the models reported here we fit a single common a parameter across all conditions. We allowed for the uninformed guessing rate to differ between single (or dual) probe and whole display ($\{u_{\text{sp}}, u_{\text{wd}}\}$). The models differ in how the allow k to vary.

References

- Adam, K. C., Vogel, E. K., & Awh, E. (2017). Clear evidence for item limits in visual working memory. *Cognitive Psychology*, *97*, 79–97.
- Allen, R. J., Baddeley, A. D., & Hitch, G. J. (2006). Is the binding of visual features in working memory resource-demanding? *Journal of Experimental Psychology: General*, *135*(2), 298–313.
- Allen, R. J., Castellà, J., Ueno, T., Hitch, G. J., & Baddeley, A. D. (2015). What does visual suffix interference tell us about spatial location in working memory? *Memory & Cognition*, *43*(1), 133–142.
- Allen, R. J., Hitch, G. J., Mate, J., & Baddeley, A. D. (2012). Feature binding and attention in working memory: A resolution of previous contradictory findings. *The Quarterly Journal of Experimental Psychology*, *65*(12), 2369–2383.
- Alvarez, G. A., & Thompson, T. W. (2009). Overwriting and rebinding: Why feature-switch

- detection tasks underestimate the binding capacity of visual working memory. *Visual Cognition*, 17(1-2), 141–159.
- Aust, F., & Barth, M. (2018). *papaja: Create APA manuscripts with R Markdown*. Retrieved from <https://github.com/crsh/papaja>
- Baddeley, A. D., Allen, R. J., & Hitch, G. J. (2011). Binding in visual working memory: The role of the episodic buffer. *Neuropsychologia*, 49(6), 1393–1400.
- Bates, D., Mächler, M., Bolker, B., & Walker, S. (2015). Fitting linear mixed-effects models using lme4. *Journal of Statistical Software*, 67(1), 1–48. doi:10.18637/jss.v067.i01
- Bays, P. M., Catalao, R. F., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10), 7.
- Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L., & Husain, M. (2011). Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*, 11(10), 6–6.
- Brady, T. F., & Alvarez, G. A. (2011). Hierarchical encoding in visual working memory: Ensemble statistics bias memory for individual items. *Psychological Science*, 22(3), 384–392.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2009). Compression in visual working memory: Using statistical regularities to form more efficient memory representations. *Journal of Experimental Psychology: General*, 138(4), 487–502.
- Brady, T. F., Konkle, T., & Alvarez, G. A. (2011). A review of visual memory capacity: Beyond individual items and toward structured representations. *Journal of Vision*, 11(5), 4.
- Chen, Z., & Cowan, N. (2013). Working memory inefficiency: Minimal information is

- utilized in visual recognition tasks. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(5), 1449–1462.
- Cowan, N. (2017). The many faces of working memory and short-term storage. *Psychonomic Bulletin & Review*, 24(4), 1158–1170.
- Cowan, N., Blume, C. L., & Saults, J. S. (2013). Attention to attributes and objects in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39(3), 731–747.
- Cowan, N., Elliott, E. M., Saults, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., & Conway, A. R. A. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, 51(1), 42–100.
- Cowan, N., Saults, J. S., & Blume, C. L. (2014). Central and peripheral components of working memory storage. *Journal of Experimental Psychology: General*, 143(5), 1806–1836.
- Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, 120(4), 873–902.
- Fox, J., & Weisberg, S. (2011). *An R companion to applied regression* (Second.). Thousand Oaks CA: Sage. Retrieved from <http://socserv.socsci.mcmaster.ca/jfox/Books/Companion>
- Gilchrist, A. L., & Cowan, N. (2014). A two-stage search of visual working memory: Investigating speed in the change-detection paradigm. *Attention, Perception, & Psychophysics*, 76(7), 2031–2050.
- Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal

- representations. *Journal of Cognitive Neuroscience*, 15(8), 1176–1194.
- Hardman, K. O., & Cowan, N. (2015). Remembering complex objects in visual working memory: Do capacity limits restrict objects or features? *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41(2), 325–347.
- Hardman, K. O., & Cowan, N. (2016). Reasoning and memory: People make varied use of the information available in working memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 42(5), 700–722.
- Hu, Y., Hitch, G. J., Baddeley, A. D., Zhang, M., & Allen, R. J. (2014). Executive and perceptual attention play different roles in visual working memory: Evidence from suffix and strategy effects. *Journal of Experimental Psychology: Human Perception and Performance*, 40(4), 1665.
- Hyun, J.-S., Woodman, G. F., Vogel, E. K., Hollingworth, A., & Luck, S. J. (2009). The comparison of visual working memory representations with perceptual inputs. *Journal of Experimental Psychology: Human Perception and Performance*, 35(4), 1140–1160.
- Jiang, Y., Olson, I. R., & Chun, M. M. (2000). Organization of visual short-term memory. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 26(3), 683–702.
- Johnson, J. S., Hollingworth, A., & Luck, S. J. (2008). The role of attention in the maintenance of feature bindings in visual short-term memory. *Journal of Experimental Psychology: Human Perception and Performance*, 34(1), 41–55.
- Kondo, A., & Saiki, J. (2012). Feature-specific encoding flexibility in visual working memory. *PloS ONE*, 7(12), e50962.
- Liesefeld, H. R., Liesefeld, A. M., & Müller, H. J. (2019). Two good reasons to say

- “change!”—ensemble representations as well as item representations impact standard measures of vwm capacity. *British Journal of Psychology*, 110(2), 328–356.
- Logie, R. H. (2011). The functional organization and capacity limits of working memory. *Current Directions in Psychological Science*, 20(4), 240–245.
- Logie, R. H., Brockmole, J. R., & Jaswal, S. (2011). Feature binding in visual short-term memory is unaffected by task-irrelevant changes of location, shape, and color. *Memory & Cognition*, 39(1), 24–36.
- Logie, R. H., Brockmole, J. R., & Vandenbroucke, A. R. (2009). Bound feature combinations in visual short-term memory are fragile but influence long-term learning. *Visual Cognition*, 17(1-2), 160–179.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Macmillan, N. A., & Creelman, C. D. (2005). *Detection theory: A user's guide* (2nd ed.). Mahwah, New Jersey: Lawrence Erlbaum Associates.
- Makovski, T., Watson, L. M., Koutstaal, W., & Jiang, Y. (2010). Method matters: Systematic effects of testing procedure on visual working memory sensitivity. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 36(6), 1466–1479.
- Morey, C. C., & Bieler, M. (2013). Visual short-term memory always requires general attention. *Psychonomic Bulletin and Review*, 20, 163–170.
- Oberauer, K., & Eichenberger, S. (2013). Visual working memory declines when more features must be remembered for each object. *Memory & Cognition*, 41(8), 1212–1227.
- Pashler, H. (1988). Familiarity and visual change detection. *Perception and Psychophysics*,

44(4), 369–378.

Phillips, W., & Christie, D. (1977). Components of visual memory. *The Quarterly Journal of Experimental Psychology*, 29(1), 117–133.

R Core Team. (2018). *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. Retrieved from <https://www.R-project.org/>

Revelle, W. (2018). *Psych: Procedures for psychological, psychometric, and personality research*. Evanston, Illinois: Northwestern University. Retrieved from <https://CRAN.R-project.org/package=psych>

Rhodes, S., & Cowan, N. (2019). Flexible representations in visual working memory and interactions with long-term learning: Commentary on the special issue. *British Journal of Psychology*, 110(2), 449–460.

Rhodes, S., Cowan, N., Hardman, K. O., & Logie, R. H. (2018). Informed guessing in change detection. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 44(7), 1023–1035.

Rhodes, S., Cowan, N., Parra, M. A., & Logie, R. H. (2018). Interaction effects on common measures of sensitivity: Choice of measure, type i error, and power. *Behavior Research Methods*, online ahead of print. doi:10.3758/s13428-018-1081-0

Rhodes, S., Parra, M. A., & Logie, R. H. (2016). Ageing and feature binding in visual working memory: The role of presentation time. *The Quarterly Journal of Experimental Psychology*, 69(4), 654–668.

Rhodes, S., Parra, M. A., Cowan, N., & Logie, R. H. (2017). Healthy aging and visual working memory: The effect of mixing feature and conjunction changes. *Psychology*

- and Aging*, 32(4), 354–366.
- Rouder, J. N., Morey, R. D., Cowan, N., Zwillling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of fixed-capacity models of visual working memory. *Proceedings of the National Academy of Sciences*, 105(16), 5975–5979.
- Rouder, J. N., Morey, R. D., Morey, C. C., & Cowan, N. (2011). How to measure working memory capacity in the change detection paradigm. *Psychonomic Bulletin and Review*, 18, 324–330.
- Schneegans, S., & Bays, P. M. (2019). New perspectives on binding in visual working memory. *British Journal of Psychology*, 110(2), 207–244.
- Schneider, W., Eschman, A., & Zuccolotto, A. (2012). *E-prime user's guide*. Pittsburgh: Psychology Software Tools, Inc.
- Shimi, A., & Logie, R. H. (2019). Feature binding in short-term memory and long-term learning. *Quarterly Journal of Experimental Psychology*, 72(6), online ahead of print.
- Tabi, Y. A., Husain, M., & Manohar, S. G. (2019). Recall cues interfere with retrieval from visuospatial working memory. *British Journal of Psychology*, 110(2), 288–305.
- Treisman, A. (1977). Focused attention in the perception and retrieval of multidimensional stimuli. *Perception & Psychophysics*, 22(1), 1–11.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12(1), 97–136.
- Ueno, T., Allen, R. J., Baddeley, A. D., Hitch, G. J., & Saito, S. (2011a). Disruption of visual feature binding in working memory. *Memory and Cognition*, 39, 12–23.
- Ueno, T., Mate, J., Allen, R. J., Hitch, G. J., & Baddeley, A. D. (2011b). What goes

- through the gate? Exploring interference with visual feature binding. *Neuropsychologia*, 49, 1597–1604.
- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, 109(22), 8780–8785.
- 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(1), 92–114.
- Wheeler, M. E., & Treisman, A. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131(1), 48–64.
- Wickham, H., & Henry, L. (2018). *Tidyr: Easily tidy data with 'spread()' and 'gather()' functions*. Retrieved from <https://CRAN.R-project.org/package=tidyr>
- Wolfe, J. M., & Cave, K. R. (1999). The psychophysical evidence for a binding problem in human vision. *Neuron*, 24(1), 11–17.
- Woodman, G. F., Vogel, E. K., & Luck, S. J. (2012). Flexibility in visual working memory: Accurate change detection in the face of irrelevant variations in position. *Visual Cognition*, 20(1), 1–28.
- Yang, C.-T., Tseng, P., & Wu, Y.-J. (2015). Effect of decision load on whole-display superiority in change detection. *Attention, Perception, & Psychophysics*, 77(3), 749–758.
- Yeh, Y.-Y., Yang, C.-T., & Chiu, Y.-C. (2005). Binding or prioritization: The role of selective attention in visual short-term memory. *Visual Cognition*, 12, 759–799.
- Zhang, H., Xuan, Y., Fu, X., & Pylyshyn, Z. W. (2010). Do objects in working memory

compete with objects in perception? *Visual Cognition*, 18(4), 617–640.