Location-based errors in change detection: A challenge for the slots model of visual working memory

Chris Donkin · Sophia Chi Tran · Mike Le Pelley

Published online: 18 November 2014
© Psychonomic Society, Inc. 2014

Abstract The slots model of visual working memory, despite its simplicity, has provided an excellent account of data across a number of change detection experiments. In the current research, we provide a new test of the slots model by investigating its ability to account for the increased prevalence of errors when there is a potential for confusion about the location in which items are presented during study. We assume that such location errors in the slots model occur when the feature information for an item in one location is swapped with the feature information for an item in another location. We show that such a model predicts two factors that will influence the extent to which location errors occur: (1) whether the test item changes to an “external” item not presented at study, or to an “internal” item presented at another location during study, and (2) the number of items in the study array. We manipulate these factors in an experiment, and show that the slots model with location errors fails to provide a satisfactory account of the observed data.

Keywords Working memory · Short term memory · Memory · Change detection · Memory models

Introduction

In the past decade, much research has been dedicated to developing a model of how information is stored in visual working memory (VWM). For the purposes of this paper, the term VWM refers to the active maintenance of visual information in mind for up to several seconds. At present, existing theories of the capacity of VWM can be roughly divided into

two opposing classes: (a) slots models, and (b) resource models. According to the approach taken by slots models, VWM consists of a fixed number of discrete memory slots, usually between three and five, each capable of storing one whole visual item with high precision (Cowan, 2001; Luck & Vogel, 1997; Rouder et al., 2008; Zhang & Luck, 2008). In contrast, the resource models propose that VWM uses a limited pool of mnemonic resource that can be flexibly distributed across a variable number of items (Bays, Catalao & Husain, 2009; Wilken & Ma, 2004; van den Berg, Shin, Chou, George & Ma, 2012).

The change detection task is one of the most common paradigms used to investigate VWM (Cowan et al., 2005). In this task the observer is presented with an array of visual items to study, such as colored squares. After a short retention interval, a single test square is presented at a particular location and the observer has to decide whether the test square is the same color as the square in the corresponding location of the study array.

Rouder et al. (2008) showed that a fixed-capacity slots model provided an excellent account for data in the change detection task. Compared to resource models, the slots model provided a more parsimonious account of behavior. Donkin, Tran and Nosofsky (2014) extended these results, and showed that the slots model was preferred over resource models in three out of four experiments using the change detection task. Donkin, Nosofsky, Gold and Shiffrin (2013) also found evidence for a mixture of memory-based and guessing responses when analyzing full response-time distributions, consistent with the slots model (but see van den Berg et al., 2012 for evidence of resource models that predict guessing-like behavior).

The general preference for the slots model over the resource model in the change detection task comes largely because of its simplicity. In all of the change detection studies for which the slots model is preferred, the number of items to

C. Donkin (✉) · S. C. Tran · M. Le Pelley
School of Psychology, University of New South Wales, Sydney,
NSW, Australia
e-mail: christopher.donkin@gmail.com

be remembered (or set size) was manipulated. The slots model provides a strong quantitative prediction regarding the influence of set size on performance, with the probability of a correct response being determined primarily by the probability that any given item is encoded into a slot in memory. Given this constraint, the quantitative match between the decrease in observed performance with set size and that predicted by the model is impressive. The current research provides a further test of the ability of the slots model to explain behavior in the change detection task. In particular, we investigated the slots model's account for location errors.

It is generally assumed that errors in the change detection task arise because of “feature errors,” such as when an observer fails to remember the color of a studied item (Zhang & Luck, 2008). However, Bays, Catalao and Husain (2009; see also Wheeler & Triesman, 2002) showed that participants make incorrect responses due to a combination of feature errors and “location errors”. Using a color reproduction/recall task, they showed that when asked about the identity of an item presented in a particular location at study, the most common error participants made was to report a color that was presented at study but in a different location. That is, instead of recalling the color of the item presented in the target location, participants would report the color of an item presented elsewhere in the study array. A number of subsequent studies have also found evidence for the importance of location memory in the color reproduction paradigm (Bays, Wu & Husain, 2011; Emrich & Ferber, 2012; Rajsic & Wilson, 2012).

Our aim in the present study was to examine location errors in the change detection task by manipulating the extent to which location errors were possible. To do this, we used two types of change trials. In “external” change trials, the test item would change to a color that had not been presented anywhere in the study array for that trial. In “internal” change trials, the test item would change to a color that had been presented at another location in the study array. The rationale behind this manipulation is that in the external change condition, location errors are impossible, since the change is such that the test item is different from all items presented in the study array. However, on internal change trials, participants may make a location error and thus incorrectly identify the test item as having been presented in the target location in the study array, resulting in an incorrect “same” response.

We will now summarize the slots model proposed by Rouder et al. (2008). By virtue of its simplicity, this model is easily extended to predict the occurrence of location errors in the change detection task. We then go on to test whether or not these predictions are consistent with human data.

The slots model

Rouder et al.'s (2008) fixed-capacity slots model assumes that observers store a fixed number of k items in VWM with high precision. In the change detection task, N items are presented for study, and one of those items is then probed at test. The probability that the test item is one of the k items stored in memory, denoted d , is $\min(1, \frac{k}{N})$. For example, if a participant with a capacity of four items is presented with six items to study, the probability that the test item is stored in memory is $d = \frac{4}{6}$. If only two items are presented to this participant, then the probability that the test item is in memory will instead be $d = 1$.

If the target item is stored in memory, it is assumed that the participant will respond correctly. However, for trials on which the target item was not stored, the participant has no information about that item, and so will guess “change” with probability g . The g parameter is such that a value of 0.8 indicates that participants will respond that the test item changed on 80 % of trials for which they have no memory about the test item.

As described, the slots model predicts perfect performance when set size is smaller than capacity (since in this case we have $d = 1$). However, participants generally make some number of incorrect responses on all trials, regardless of difficulty. In order to allow the slots model to predict errors for small set sizes, Rouder et al. (2008) assumed that observers pay attention with probability a . When observers are attentive, performance is as previously described. However, when observers fail to pay attention on a trial, they have no memory regarding the target item, and so guess “change” with probability g .

Based on these assumptions, the hit rate (i.e., correct change detections) and false alarm rate (i.e., incorrect change detections) are given by

$$h = a(d + (1-d)g) + (1-a)g \quad (1)$$

$$f = a(1-d)g + (1-a)g \quad (2)$$

Incorporating location errors

Location errors involve the observer confusing the locations in which items were presented. In the slots model, we assume that when a particular location is probed with a test item, observers will first determine whether they have a memory of an item bound to that location (cf. Kahneman, Triesman & Gibbs, 1992). If an item is present, they will retrieve the features of the item for that location (if there is no memory of an item in that location, the observer will guess). We assume that a location error occurs when the feature information for an item in one location is swapped with the feature information for an item in another location.

Figure 1 gives an example of a location error, and the implications for performance. In the example, a study array contains a red item in location L1, a green item in location L2, and an orange item in location L3. Figure 1A represents a perfect encoding of the study array, and so whenever the item in location L1 is probed at test, the feature “red” will be correctly retrieved and the response will be correct. Figure 1B represents a location error, such that the feature “red” is incorrectly bound to location L2 and the feature “green” is bound to location L1. Now, when the participant is probed with a test item at location L1, their memory will indicate that a green item was present at study, and this may lead to incorrect responses.

The first test array of Fig. 1B shows that location errors will always lead to an incorrect response on same trials (i.e., false alarms). When the test item is a red item presented in location L1 (i.e., the test item is the same as the study item), because the observer’s memory is that the item in location L1 was green, they will make an incorrect “change” response. Let B be the probability of a correct location-feature binding of the test item (so that the probability of a location error is given by $[1-B]$), then the false alarm rate in the slots model now becomes

$$f = a(d(1-B) + (1-d)g) + (1-a)g \quad (3)$$

The second example of a test array in the rightmost column shows that location errors have no impact on external change trials. In the figure, the test item in location L1 is yellow (i.e., the test item changes to a color not present in the study array).

Despite the location error, the test item is still different to the observer’s memory for that item, and so a correct “change” response will be made. As such, the hit rate for external changes, h_{EXT} , will be equal to h as given by Eq. 1.

The third and fourth test arrays of Fig. 1B show that location errors on internal change trials will sometimes lead to incorrect “same” responses. In the third test array, the test item in location L1 changes from red to the green item originally presented in location L2. Because the observer made a location error, their memory indicates that a green item was presented in location L1. As such, the participant incorrectly responds that the study and test items are the same. However, for the fourth test array, when the test item is the item from the study array for which a location error did not occur, then a correct “change” response is still made.

The likelihood with which location errors will lead to a miss for test arrays containing an internal change will be determined largely by set size. To see this, consider a study array that now contains six items, and which is encoded with a location error. Test arrays featuring an internal change still give two possibilities. If the test item is the same as the item for which a location error was made, then the participant will make an incorrect “same” response. However, there are now four other colors in the display to which the test item could potentially change that were not involved in the location error. As such, the probability that a location error will lead to an incorrect “same” response has decreased. More formally, the probability that a location error will lead to an incorrect “same” response is $\frac{1}{N-1}$. As such, the probability of a correct “change” response, even though a location error has been

Fig. 1 Graphical representation of the predictions of the slots model for same, external, and two types of internal change trials, without a location error (left) and with a location error (right)

made, is given by $\frac{N-2}{N-1}$. The proportion of hit responses for test arrays involving an internal change, h_{INT} , is therefore given by

$$h_{INT} = a \left(d \left[B + (1-B) \frac{N-2}{N-1} \right] + (1-d)g \right) + (1-a)g \quad (4)$$

In the previous equations, B is the probability that the item in the probed location did not suffer a location error. This probability might be influenced by a number of factors. We present the results of one possibility, but detail alternatives in an [Appendix](#). In the model we report here, we assume that as more items are presented, the probability of a location error increases. The idea is that each item in the study array has a chance of being switched with the item in the probed test location. As such, if b is the probability that the item in the probed location is not switched with another particular item, the probability of none of the N items being switched with the item in the probed test location, B , is b^{N-1} .

However, such a model assumes that all items have the potential to cause a location error. The fundamental concept behind the slots model is that not all items are encoded into memory, and therefore may not have an opportunity to cause a location error. As such, we further assume that only items that are in memory can cause a location error. The probability of no location errors for the item in the probed location now becomes

$$B = b^{\min(k-1, N-1)} \quad (5)$$

Qualitative predictions

We report an experiment directed at testing two qualitative predictions made by this modified slots model. The first prediction is that the hit rate in the internal change condition, h_{INT} , should be lower than that in the external change condition, h_{EXT} , due to location errors. Consider Eq. 4 above. Since $0 \leq B \leq 1$, then $(B + (1-B) \frac{N-2}{N-1}) \leq 1$. Hence comparing Eqs. 1 and 4, we have $h_{EXT} \geq h_{INT}$ (where $h_{EXT} = h_{INT}$ only if the probability of location errors is zero). The second prediction is that the difference between the hit rate in the internal and external change conditions should become smaller as set size becomes larger. This is because $(B + (1-B) \frac{N-2}{N-1}) \rightarrow 1$ as $N \rightarrow \infty$, so comparing Eqs. 1 and 4 gives $h_{EXT} \rightarrow h_{INT}$ as $N \rightarrow \infty$.

Quantitative predictions using prior predictives

One troubling aspect of our second qualitative prediction is that all reasonable models of working memory will predict that $h_{EXT} \rightarrow h_{INT} \rightarrow f$ as $N \rightarrow \infty$. That is, unless a model assumes unlimited memory capacity, performance must decrease to chance as the number of items to remember approaches infinity. However, we can show that the slots model predicts that $h_{EXT} \rightarrow h_{INT}$ as set size increases over a range that is commonly used in visual working memory studies.

The left panel of Fig. 2 plots the prior predictives for the difference between h_{EXT} and h_{INT} as a function of set size. Prior predictives are the predictions from the model that arise based on the prior distributions of the parameters of the model. In other words, Fig. 2 shows the most likely values of $h_{EXT} - h_{INT}$ for each set size, as predicted by the slots model. Prior predictives are a useful tool for investigating the predictions of a model, as prior distributions form an integral part of any theory, representing what we know about the parameters of the model (Vanpaemel, 2010; Vanpaemel & Lee, 2011).

The prior distributions of the main parameters in the model (a , k , g , and b) are shown in the right columns of Fig. 2 (see Table 2 for the full definition of all prior distributions). The prior distributions we used were based on the data from Experiment 1 in Donkin, Tran & Nosofsky (2014). Their experiment was a replication of Rouder et al. (2008), and included a between-subjects manipulation of external and internal change, making it possible to calculate an approximate estimate of the probability of location errors, b . To construct the prior distributions in Fig. 2 we added additional variability to the posterior distributions estimated from Donkin et al.'s experiment, so as to highlight the robustness of the slots model's prediction.

The open circles in the leftmost column of Fig. 2 represent the median value of the posterior difference between h_{EXT} and h_{INT} , which shows that the slots model predicts that the difference between the two hit rates is almost zero when $N=8$. The error bars represent the central 95 % of the prior predictive distribution of $h_{EXT} - h_{INT}$, and show that the slots model clearly predicts a decrease in $h_{EXT} - h_{INT}$ as N increases from 2 to 8. In other words, based on our prior predictive analysis, we expect to see a decrease in the difference between h_{EXT} and h_{INT} as set size increases from 2 to 8 items.

Method

Participants

Participants were 32 first-year psychology students from the University of New South Wales. They received course credit in exchange for one hour of participation.

Materials, stimuli and design

The experiment was conducted on a standard desktop computer and responses were made on a keyboard. All instructions and stimuli were presented on a 24-in LCD monitor (1920 × 1080 resolution).

Stimuli were color squares, 0.75×0.75 degrees of visual angle in size, at a viewing distance of 60 cm. The color of each square was randomly sampled without replacement from a set of ten highly discriminable colors: white, black, red, blue,

Fig. 2 The central 95 % of the prior predictive distribution for the difference between h_{EXT} and h_{INT} are plotted in the leftmost column. Prior distributions of the four parameters (a , k , g , and b) that are used to

generate these prior predictives are plotted to the right of the figure. In each panel, open circles show median values

green, yellow, orange, cyan, purple, and dark-blue-green. The study array was presented in the way described by Cowan et al. (2005). That is, squares were presented on a grey background, within a 9.8×7.3 degree rectangular array. The positioning of squares was random, but with the restriction that each square must be at least 2° away from the center of the array and from any other square.

The test array contained a single color square, presented at a randomly chosen location of the study array. A black circle with 1 pixel line-width and 1.5° diameter surrounded the test square to help cue its location. For each participant and within each block of 80 trials, the test square was the same color as the square in the corresponding location of the study array on half of the trials (i.e., a “same” trial), and changed from study to test on the other half of trials (i.e., a “change” trial). Half of the change trials had an external change, where the test item changed to a color not presented at study; this color was chosen randomly from the set of ten colors in the stimulus set, excluding any that had been presented in the study array. The other half were internal change trials, where the test item changed to a color presented at another location in the study array, with this location being randomly chosen.

The study array contained two, four, six, or eight color squares. The number of items, or set size, varied randomly

from trial-to-trial, with the restriction that each set size was presented an equal number of times within a given block of trials. Participants completed six blocks of 80 trials, yielding a total of 60 same trials, 30 external change trials and 30 internal change trials for each set size condition. Each participant first completed four practice trials. They were encouraged to take a self-paced rest period between each block of trials. The experiment took approximately 40 minutes to complete.

Procedure

At the beginning of each trial, a fixation cross appeared in the center of the screen for 1000 ms. A study array containing N color squares was then presented for 500 ms. This was followed by a blank screen for 500 ms, then a multicolored pattern mask consisting of all ten colors in the stimulus set was presented at each of the study square locations for 500 ms. Next, a test array was presented. Using the keyboard, participants were required to indicate whether the color of the square in the test array was the same as the corresponding study square, or had changed, by pressing the “J” or “F” key, respectively. After the participant made their response, response feedback was displayed on the screen for 1000 ms. After a 1000-ms blank screen, the next trial began.

Statistical approach

In the current article we use Bayesian methods to analyze our data. For our standard analyses, we use the default Bayesian ANOVAs outlined by Rouder, Morey, Speckman & Province (2012; see also Wetzels, Grasman, & Wagenmakers, 2012). Doing so avoids the well documented disadvantages of null-hypothesis hypothesis testing (e.g., Wetzels et al., 2011). The Bayes factors from Bayesian ANOVAs give the relative likelihood of the data under the null and alternative hypothesis (for additional discussion, see Morey, Rouder, Verhagen & Wagenmakers, 2014; Wagenmakers, Lodewyckx, Kuriyal, & Grasman, 2010). Bayes factors not only allow us to indicate whether the data support the null or the alternative, but also quantify the strength of the support for one hypothesis over the other.

We also use Bayesian estimation methods to obtain posterior distributions of the parameters of the modified slots model. Analyzing posterior distributions, instead of best-fitting parameters, allows us to measure the uncertainty in our parameter estimates. Bayesian methods also easily permit hierarchical modeling, in which we infer values of parameters for populations at the same time as for individuals (Lee, 2011). As argued by Morey (2011), hierarchical models are essential for estimating visual working memory capacity in change detection tasks, as they correct a number of serious problems inherent to non-hierarchical methods. The model we fit is a modified version of that offered in Morey and Morey (2011).

Results

Trials with responses faster than 200 ms or slower than 3 s were removed. A total of 1.1 % of the data was censored.

The unfilled points in Fig. 3 show the proportion of trials on which the participant made a change response, as a function of set size, for the external change, internal change, and same conditions. For the external change and internal change conditions, a high proportion of change responses is indicative of good performance (i.e., high hit rate), while for the same condition, a high proportion of change responses indicates poor performance (i.e., high false alarm rate). As is clear from the figure, performance declines as the number of items in the study array becomes larger. A Bayesian one-way, repeated-measures ANOVA (Set Size: 2,4,6,8) on the proportion of correct responses yielded a Bayes Factor of 6.9×10^{36} , suggesting the data were much more likely under the alternative hypothesis of an effect of set size than the null hypothesis of no effect of set size.¹

¹ We performed the common arcsine transformation of the square-root of proportions in order to satisfy equal-variance assumptions of the ANOVA.

Fig. 3 The proportion of trials on which the participant responded “change”, averaged across individuals, as a function of set size, for each trial type; external change, internal change, and same (red, black, and green, respectively). The model fits of the slots model are overlaid in the form of posterior predictives. The external and internal conditions are offset horizontally for clarity

Figure 3 also shows that, as predicted by the modified slots model, there is a higher hit rate in the external change condition than in the internal change condition. However, unlike the qualitative predictions from the slots model, the difference between the internal and external change conditions does not appear to get smaller as set size becomes larger. Indeed, if anything the difference in hit rate between internal and external change conditions becomes larger with increasing set size.

We conducted a 2 (Change: internal, external) $\times$ 4 (Set Size: 2,4,6,8) repeated-measures Bayesian ANOVA on the proportion of “change” responses. Table 1 contains Bayes factors for the ANOVA analysis. The Bayes factor is largest for the model with only the main effects of Change and Set Size, being 1.5 times more likely than the model that also includes an interaction. That is, though the difference between hit rates in external and internal change conditions appears to increase with set size (see Fig. 3, and also Fig. 5), it is unclear whether this interaction is statistically reliable. The largest effect was that of Set Size (more likely than a model without a main effect of Set Size by a factor of 10^{21}), though the effect of Change is also highly reliable (by a factor of 10^{14}).²

² Bayes factors are defined by the relative likelihood of the data under the two models. However, since the models are wrong, these probabilities are unlikely to be absolutely true. It is worth noting that the same objection is true of the frequentist ANOVA, and the resultant p-values. Thankfully, the effects of set size and change type are clear enough in our data that such assumptions are unlikely to be critical to our interpretation of the data.

Table 1 Bayes factors for 2 (Change: internal, external) $\times$ 4 (Set Size: 2,4,6,8) repeated-measures Bayesian ANOVA on the proportion of “change” responses. Bayes factors are relative probabilities against the model with no factor included (i.e., only the random effect of Subject). The natural logarithms of the Bayes factors are also reported

Model	Bayes factor	ln(Bayes Factor)
Change + Subject	1.43×10^5	9.57
Set size + Subject	4.35×10^{19}	45.22
Change + Set Size + Subject	9.15×10^{26}	62.08
Change + Set Size + Change*Set Size + Subject	6.22×10^{26}	61.69

The overall pattern of data seems inconsistent with the qualitative predictions of the slots model. However, since our intuition for the behavior of models is relatively poor, we also fitted the model to the data to examine whether any combination of parameter values allowed it to explain our results.

Fitting the model

As in Rouder et al. (2008), each participant’s hit and false alarm responses from external change, internal change, and same trials were assumed to have come from a binomial distribution. The rates of these binomial distributions were generated using Eqs. 1 to 5 for f , h_{EXT} , and h_{INT} . These equations require four free parameters, k , g , a , and b , none of which vary over any of the experimental conditions. As such, each individual’s 12 hit and false alarm rates were to be explained using just four free parameters.

We used Bayesian estimation methods to apply a hierarchical version of the model to our data. In particular, each individual’s parameters were assumed to have come from a hierarchical, population-level distribution. Guessing (g) and attention (a) parameters were assumed to have come from normal distributions, truncated between 0 and 1, with means μ_g and μ_a , respectively (and standard deviations σ_g and σ_a). Individual capacity estimates k were assumed to come from a population normal distribution with mean μ_k (and standard deviation σ_k) that was truncated between 1 and 8. We used a beta distribution for the population distribution for the b parameter (parameters α_b and β_b), as we expected to observe a value close to the upper bound, thus making the truncated normal distribution less ideal. The mean of the beta distribution, μ_b , is calculated as $\alpha_b / (\alpha_b + \beta_b)$.

Table 2 gives the full details of the prior distributions we used (also see Fig. 2 for plots of the priors we used for the μ parameters). Posterior distributions were obtained using Just Another Gibbs Sampler (Plummer, 2003), using seven chains of 30,000 samples, after 10,000 burn-in samples, and keeping only every 20th sample. This yields a total of 7,000 samples to make the posterior distribution for each parameter.

Table 2 Prior distributions for the parameters of the slots model. These were used to generate Fig. 2 and to estimate posterior distributions. The priors of the parameters μ_a , μ_g , α_a , and β_b followed a Beta (A, B) distribution with values of A and B given in the table. The prior of μ_k followed a normal distribution with mean (M) and standard deviation (SD) shown in the table. The priors of σ_g , σ_a and σ_k followed a uniform distribution with lower (L) and upper (U) bounds shown in the table

Parameter	Distribution	A	B
μ_a	Beta	30	5
μ_g		10	10
α_b		8	160
β_b		1	15
		M	SD
μ_k	Normal	3.5	1.5
		L	U
$\sigma_g, \sigma_a, \sigma_k$	Uniform	0.01	10

We first look at the estimated parameters of the slots model. The top row of Fig. 4 plots the posterior distributions for the means of the hierarchical distributions for μ_a , μ_g , μ_k , and μ_b . These posterior distributions tell us what values of the population means are most likely, given the observed data. Participants appeared to have paid attention on most trials, as μ_a is close to a ceiling value of 1. The average capacity estimate for the population, μ_k , is most likely around 2.5 to 3, which is slightly lower than is usually observed (Cowan, 2001; Donkin, Tran & Nosofsky, 2014; Rouder et al., 2008).³ The posterior distribution of the average guessing rate, μ_g , shows that participants appeared to have been biased to guess “change” with a probability of approximately 0.65.

The population average probability of two items in memory not being switched, μ_b , is around 0.96. This suggests that the model predicts a relatively small chance of a location error between any two items. That said, the posterior distribution for μ_b has almost no mass at 1, which suggests that the model does infer the existence of location errors. The second row of Fig. 4 plots the posterior probability of no location error, B , for each set size. The probability of no location error drops as set size increases from 2 to 4, but is relatively stable for larger set sizes. This is because most participants’ capacity did not exceed four items, therefore restricting the number of potential location errors (since $B = b^{\min(k-1, N-1)}$). We now examine how well this model can account for the precise pattern of hit rates across external and internal change conditions.

³ Note that capacity is not estimated directly from hit and false alarm rates by Pashler (1988) or Cowan (2001)’s formulae. Instead they are inferred using equations y1, 3, 4, and 5 in our Hierarchical Bayesian model. Our estimate of capacity is most similar to the approach taken by Morey (2011) and Rouder et al. (2008).

Fig. 4 The first row plots the posterior distributions of the means of the hierarchical population-level distributions of the four parameters of the slots model: μ_a , μ_k , μ_g , and μ_b , respectively. The second row plots the

posterior distribution of the probability of no location error for the probed item, B , for each of the four set size conditions

Figure 3 plots the posterior predictive model fits as crosses. The range indicated by the error bars represents the central 95 % of posterior samples of predicted hit and false alarm rates for the entire set of participants. These posterior predictives incorporate uncertainty in the parameter estimates, as well as the variability across individuals' hit and false alarm rates under the hierarchical model. The model provides a good account of false alarm rates, and it also accounts for the difference in hit rates between internal and external change conditions for a set size of 2. However, the model predicts too small a difference between the two change conditions for the larger set sizes.

Figure 5 highlights the inadequacy of the model. The plot shows the observed difference in hit rates between external and internal change conditions, averaged across individuals. The overlaid posterior predictives for the model show the predicted difference in hit rates. It is clear that as set size becomes larger, the model predicts increasingly smaller differences between the external and internal change conditions. However, the data suggest that, if anything, the difference is larger for set sizes greater than 2 (though it is unclear whether this effect is reliable).

Discussion

Our data clearly indicate that memory for the location in which items are presented is important in the change detection task. When the test item changed to an item presented elsewhere in the study array, errors increased by 10 % relative to when the probed item changed to an item that was not present at study. This result is consistent with prior research using the color reproduction/recall task, in which participants incorrectly recall items presented in non-target locations (Bays et al., 2009; Bays et al. 2011; Emerich & Ferber, 2012; Rajsic & Wilson, 2012; Wheeler & Triesman, 2002).

We modified the slots model for change detection to account for errors in the binding of location and feature information. The model assumes that on some trials, participants will swap the feature information bound to a particular location with the feature information bound to another location. We showed that the model predicts that such errors in feature-location binding will have a smaller influence as the number of items in the study array increases. Contrary to these predictions, the observed data indicate that the difference in hit rates between external and internal change conditions remains

Fig. 5 The difference in hit rates between internal and external change conditions, averaged across individuals, for each set size condition (solid squares). Error bars are standard errors of the mean. Model predictions are overlaid as the central 95 % of the posterior predictives (open circles show median values)

consistent across the set sizes or may increase with set size, at least across the range of set sizes (2–8 items) that we studied.

Our results are broadly consistent with those reported by Wheeler and Treisman (2002). They too demonstrated that performance was worse when items underwent the equivalent of our internal change manipulation, and argued against Luck and Vogel's (1997) original slots model. Most relevant to our work is their Experiment 3B, which included the equivalent of our “internal” and “external” change conditions, called “binding” and “color” conditions, respectively. Interestingly, and unlike our results, they found no difference between these two conditions, even as set size increased from three to six items. A key distinction between these experiments was that Wheeler and Treisman's participants were presented with the internal and external change conditions in blocks of trials, and were explicitly informed about the type of change they were to observe. In our experiment, unbeknownst to participants, either an internal or external change could occur on any given trial. That our participants were worse for internal change and Wheeler and Treisman's were not suggests that participants may be able to adjust their attention to specific features of items (i.e., color or location), depending on the demands of the task.

These data are also challenging for certain variants of the resource model of the change detection task. Any model in which the location of the test item cues the retrieval of feature information, and these features can be incorrectly stored, will struggle to account for these data. The misfit of the slots model

occurs because of the decrease in probability that any location-feature binding error will lead to an incorrect response as set size, N , increases (i.e., the $1/N$ chance that the test item is that for which a location error has occurred). So, for example, a resource-based model that assumes study items are encoded with varying precision would also fail to fit this data if it assumes that location is used to retrieve the feature information of the study item, and this feature information can be incorrectly bound to location.

Our results may instead suggest the need for an alternative explanation of how participants carry out the change detection task. One possibility is that participants use a global-familiarity approach to identify change (Donkin & Nosofsky, 2012b; Shiffrin & Steyvers, 1997). For example, consider a model in which the presentation of a test item leads to a familiarity signal from visual working memory. If the familiarity is larger than some criterion, then a “same” response is elicited. However, if the familiarity is weaker than the criterion, a “change” response is given. If we also assume that the strength of the familiarity signal from memory is based on both feature and location information, then such a model would predict a difference between internal and external change conditions.

Interestingly, the signal from a global-familiarity model could be driven by memory that is based on either slots or a continuous resource (see Donkin & Nosofsky, 2012a, for one such example). The model proposed by Keshvari, van den Berg and Ma (2012) seems like a good start to such an endeavor. Their model incorporates uncertainty in the features of stimuli when performing a change detection task, and could presumably be extended to also incorporate uncertainty in the location of items. If participants take advantage of both sources of uncertainty when making their decisions, then it may be possible to account for our data.

An alternative explanation for the difference between the external and internal change conditions is that performance in the change detection task is not solely a function of processes occurring in visual working memory, but also relies on encoding of the stimuli in the study array via visual attention. Treisman's seminal Feature Integration Theory (FIT: Treisman, 1988; Treisman & Gelade, 1980) proposes that certain basic features of a visual array—including color—can be processed rapidly and in parallel across the visual field. Applied to the current experiment, this would suggest that participants encode the identities of the set of different colors that are present in the study array in a manner that is independent of set size. Assuming perfect memory of this set, a test item in the external change condition can then be easily identified as a change because it is rendered in a color that does not belong to this set. In contrast, a test item in the internal change condition cannot be identified as a change on this basis since it is rendered in a color that does belong to the parallel-encoded set. Identifying whether an item has

changed relies on knowing that a particular color was presented *at a particular location* in the study array, and this conjunction information is not contained in the parallel-encoded feature list of colors. Instead, it relies on participants having encoded a “bound” representation of color and location, which (according to FIT) results from a serial encoding process in which selective attention is applied to the location of each item in the study array in turn. To the extent that participants may not have time to encode all of the color-location conjunctions in the study array, FIT anticipates that performance will be worse on internal change trials than external change trials.

Clearly, FIT does not offer a full account of our data. For example, it suggests that performance in the external change condition should be independent of set size, and yet Fig. 2 shows this is not the case. The implication is that memory/capacity limitations are also important in shaping participants’ decisions. That is, performance in the change detection task may be best understood in terms of an interaction of encoding and memory processes. Future work could attempt to separate out the contribution of these two factors.

In the present study, we aimed to extend recent work on location errors in VWM tasks (e.g., Bays et al., 2009). By manipulating whether the test item changed to an internal item that had been presented in another location of the study array or an external item that was not presented at study, we successfully influenced the number of location errors that participants made. However, a slots model that incorporates location errors as a product of incorrect binding of feature and location information failed to account for the observed data.

Appendix: Alternative slots models

In this appendix we show that a number of alternative slots models, with alternative assumptions regarding location errors, also fail to fit our data. In the model presented in the body of the manuscript, we assumed that the probability of no location errors with the probed item, B , was given by $b^{\min(k-1, N-1)}$. This choice was based on the assumption that only items in memory would lead to location errors, and that items would no longer enter memory once all slots were full. An alternative model could assume that all items in the display may enter memory sequentially, but that additional items replace existing items until a set of k items remain in memory at the end of encoding. Under such assumptions, each item has an opportunity to interfere with the binding of the feature and location information of the test item. As such, the probability of no location errors occurring with the probed item would instead be $B=b^{N-1}$.

Fig. 6 The difference in hit rates between internal and external change conditions, averaged across individuals, for each set size condition (solid squares). Error bars are standard errors of the mean. Model predictions for the model in which $B=b^{N-1}$ are overlaid as the central 95 % of the posterior predictives (open circles)

Figure 6 shows that this alternative model predicts a larger difference between the internal and external change conditions for larger set sizes than the model presented in the main text (compare with Fig. 5). However, the predicted difference is still much smaller than the observed difference.

Fig. 7 Same format as Figure A1, except that predictions are shown for the model in which μ_b is estimated freely for each set size condition

We wanted to be certain that our assumptions about the probability of a location error occurring were not responsible for the failure of the slots model. Thus, we fit a model in which the probability of no location errors for the probed item, B , was estimated separately for each set size. As shown in Figure 7, this model also failed to provide a reasonable account of the difference between the internal and external change conditions.

To understand why even this most flexible model could not fit the data, recall that with eight items in the display, no matter the probability that the probed item experiences a location error, there is still only a $1/7$ chance that the error will lead to an incorrect response (since the probed item must be the one for which a switch occurred). To overcome such a small probability, the model must assume a large probability of a location error. However, decreasing b (and, therefore, B) also increases the predicted false alarm rate (see Eq. 3 in the main text). Now, in order to reduce the effect of an increase in the probability of a location error, the capacity parameter in the model must increase to reduce false alarms. However, now the model will fail to predict the overall effect of set size. Since the effect of set size is larger than the effect of internal and external changes, the model estimates parameters that poorly fit the difference between internal and external change conditions, but capture the effect of set size, on average.

References

- Bays, P. M., Catalao, R. F. G., & Husain, M. (2009). The precision of visual working memory is set by allocation of a shared resource. *Journal of Vision*, 9(10), 1–11. 7.
- Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia*, 49, 1622–1631.
- Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, 24, 87–185.
- Cowan, N., Elliott, E. M., Saults, S. J., Morey, C. M., 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, 42–100.
- Donkin, C., & Nosofsky, R. M. (2012a). A power-law model of psychological memory strength in short-term and long-term recognition. *Psychological Science*, 23, 625–634.
- Donkin, C., & Nosofsky, R. M. (2012b). The structure of short-term memory scanning: An investigation based on response time distributions. *Psychonomic Bulletin & Review*, 19, 363–394.
- Donkin, C., Nosofsky, R. M., Gold, J. M., & Shiffrin, R. M. (2013). Discrete-slots models of visual working-memory response times. *Psychological Review*, 120, 873–902.
- Donkin, C., Tran, S. C., & Nosofsky, R. M. (2014). Landscaping analyses of the ROC predictions of discrete-slots and signal-detection models of visual working memory. *Attention, Perception & Psychophysics*, 120(4), 873–902.
- Emrich, S. M., & Ferber, S. (2012). Competition increases binding errors in visual working memory. *Journal of Vision*, 12(4), 1–16. 12.
- Kahneman, S., Treisman, A., & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cognitive Psychology*, 24, 175–219.
- Keshvari, S., van den Berg, R., & Ma, W. J. (2012). Probabilistic computation in human perception under variability in encoding precision. *Plos One*, 7(6), e40216.
- Lee, M. D. (2011). How cognitive modeling can benefit from hierarchical Bayesian models. *Journal of Mathematical Psychology*, 55, 1–7.
- Luck, S. J., & Vogel, E. K. (1997). The capacity of visual working memory for features and conjunctions. *Nature*, 390, 279–281.
- Morey, R. D. (2011). A Bayesian hierarchical model for the measurement of working memory capacity. *Journal of Mathematical Psychology*, 55, 8–24.
- Morey, R. D., & Morey, C. C. (2011). WoMMBAT: A user interface for hierarchical Bayesian estimation of working memory capacity. *Behavior Research Methods*, 43, 1044–1065.
- Morey, R. D., Rouder, J. N., Verhagen, J., & Wagenmakers, E.-J. (2014). Why hypothesis tests are essential for psychological science: A comment on Cumming. *Psychological Science*, 25, 1289–1290.
- Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics*, 44, 369–378.
- Plummer, M. (2003). JAGS: A program for analysis of Bayesian graphical models using Gibbs sampling.
- Rajsic, J., & Wilson, D. E. (2012). Remembering where: Estimated memory for visual objects is better when retrieving location with colour. *Visual Cognition*, 20(9), 1036–1039.
- Rouder, J. N., Morey, R. D., Cowan, N., Zwilling, C. E., Morey, C. C., & Pratte, M. S. (2008). An assessment of fixed-capacity models of visual working memory. *PNAS*, 105(16), 5975–5979.
- Rouder, J. N., Morey, R. D., Speckman, P. L., & Province, J. M. (2012). Default Bayes factors for ANOVA designs. *Journal of Mathematical Psychology*, 56, 356–374.
- Shiffrin, R. M., & Steyvers, M. (1997). A model for recognition memory: REM – retrieving effectively from memory. *Psychonomic Bulletin & Review*, 4, 145–166.
- Treisman, A. (1988). Features and objects: The fourteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*, 40A, 201–237.
- Treisman, A. M., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- 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 Science*, 109, 8780–8785.
- Vanpaemel, W. (2010). Prior sensitivity in theory testing: An apologia for the Bayes factor. *Journal of Mathematical Psychology*, 54, 491–498.
- Vanpaemel, W., & Lee, M. D. (2011). Using priors to formalize theory: Optimal attention and the generalized context model. *Psychonomic Bulletin & Review*, 19, 1047–1056.
- Wagenmakers, E.-J., Lodewyckx, T., Kuriyal, H., & Grasman, R. (2010). Bayesian hypothesis testing for psychologists: A tutorial on the Savage-Dickey method. *Cognitive Psychology*, 60, 158–189.
- Wetzels, R., Grasman, R. P. P., & Wagenmakers, E.-J. (2012). A default Bayesian hypothesis test for ANOVA designs. *The American Statistician*, 66, 104–111.
- Wetzels, R., Matzke, D., Lee, M. D., Rouder, J. N., Iverson, G. K., & Wagenmakers, E.-J. (2011). Statistical evidence in experimental psychology: An empirical comparison using 855 t tests. *Perspectives on Psychological Science*, 6, 291–298.
- Wheeler, M. E., & Treisman, A. M. (2002). Binding in short-term visual memory. *Journal of Experimental Psychology: General*, 131, 48–64.
- Wilken, P., & Ma, W. J. (2004). A detection theory account of change detection. *Journal of Vision*, 4, 1120–1135.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, 453, 233–235.