Temporal dynamics of encoding, storage, and reallocation of visual working memory

Paul M. Bays — Sobell Department, UCL Institute of Neurology, Queen Square, London, UK

Nikos Gorgoraptis — UCL Institute of Cognitive Neuroscience, Queen Square, London, UK

Natalie Wee — UCL Institute of Cognitive Neuroscience, Queen Square, London, UK

Louise Marshall — Sobell Department, UCL Institute of Neurology, Queen Square, London, UK

Masud Husain — UCL Institute of Cognitive Neuroscience, Queen Square, London, UK

The process of encoding a visual scene into working memory has previously been studied using binary measures of recall. Here, we examine the temporal evolution of memory *resolution*, based on observers' ability to reproduce the orientations of objects presented in brief, masked displays. Recall precision was accurately described by the interaction of two independent constraints: an *encoding limit* that determines the maximum rate at which information can be transferred into memory and a separate *storage limit* that determines the maximum fidelity with which information can be maintained. Recall variability decreased incrementally with time, consistent with a parallel encoding process in which visual information from multiple objects accumulates simultaneously in working memory. No evidence was observed for a limit on the number of items stored. Cuing one display item with a brief flash led to rapid development of a recall advantage for that item. This advantage was short-lived if the cue was simply a salient visual event but was maintained if it indicated an object of particular relevance to the task. These cuing effects were observed even for items that had already been encoded into memory, indicating that limited memory resources can be rapidly reallocated to prioritize salient or goal-relevant information.

Introduction

Working memory (WM) refers to a short-term store for the maintenance and manipulation of information obtained from the senses (Baddeley & Hitch, 1974; Cowan, 1995; Logie, 1995; Miller, Erickson, & Desimone, 1996). While primary sensory representations are continuously over-written by new input, representations in WM are both longer lasting and more durable (Phillips, 1974; Sperling, 1960), providing a protected workspace for input to inform perceptual judgments, decision making, and action selection.

The process by which sensory input is transferred into WM is an important topic of both behavioral and neurophysiological studies (Chun & Potter, 1995; Duncan, Ward, & Shapiro, 1994; Enns & Di Lollo, 1997; Jolicoeur & Dell'Acqua, 1998; Palva, Kulashekar, Hämäläinen, & Palva, 2011). In the visual domain, the time course of transfer has been explored using a masking procedure in which a stimulus array is replaced, after a variable exposure duration, by a pattern mask (Breitmeyer, 1984). This overwrites preceding sensory input, thereby halting its encoding into visual WM. A subsequent test of recall of the array provides an estimate of how much information was transferred in the period of exposure preceding the mask (Gegenfurtner & Sperling, 1993; Shibuya & Bundesen, 1988; Vogel, Woodman, & Luck, 2006; Woodman & Vogel, 2005).

This technique has demonstrated that encoding into WM is slower when there are more elements in the array, indicating a limit on processing capacity. Studies using this procedure have estimated the encoding rate to be on the order of 20–100 ms per item. However, the correct interpretation of this figure is debated. It may be a direct reflection of a serial process, in which integrated object representations are transferred one by one into WM

(Hoffman, 1979). Alternatively, it may be an indirect measure of a capacity-limited parallel process, in which visual input is continuously encoded into WM at a rate determined by total stimulus load (Shibuya & Bundesen, 1988). It has proven difficult to distinguish between these two hypotheses, in part because previous studies were based on binary (correct/incorrect) measures of recall performance.

In contrast to this binary approach, methods for examining the *fidelity* with which visual information is stored are becoming of increasing importance in WM research (Alvarez & Cavanagh, 2004; Bays, Catalao, & Husain, 2009; Bays & Husain, 2008; Brady, Konkle, & Alvarez, 2011; Elmore et al., 2011; Fougnie, Asplund, & Marois, 2010; Palmer, 1990; Wilken & Ma, 2004; Zhang & Luck, 2008). One consequence of this new approach has been a reconsideration of the traditional concept of WM capacity as reflecting a limited number of independent memory "slots" (typically 3–4) each storing one object (Cowan, 2001; Luck & Vogel, 1997; Pashler, 1988). Newer models instead propose a unitary working memory resource that is distributed between elements of a visual scene: the more items are stored, the less precisely each can be recalled (Alvarez & Cavanagh, 2004; Bays et al., 2009; Bays & Husain, 2008; Wilken & Ma, 2004). Critically, these models allow for flexibility in allocation, such that WM resources can be preferentially directed toward a salient or behaviorally important object to enhance the resolution of its storage (Bays & Husain, 2008, 2009).

Here, we investigated the temporal evolution of working memory precision, based on observers' ability to reproduce the orientations of objects presented in masked displays of varying size and duration. We characterize two independent constraints on WM capacity: a storage limit that determines the maximum fidelity with which visual information can be maintained and an independent encoding limit that sets the rate at which this capacity is filled.

We further examined the process of memory reallocation by cuing a single item within the memory array. The results demonstrate changes in recall precision consistent with a *redistribution* of resources toward the cued item, with a corresponding cost to uncued items. The time course of reallocation depends on the behavioral relevance of the salient cue event, indicating a competition between bottom-up and top-down influences for control of the contents of WM. Recall precision provides a simple but effective index to track the deployment of working memory resources over time.

General methods

Procedure

A total of 68 subjects (25 males, 43 females, aged 18–36 years) participated in the study after giving informed consent. All subjects reported normal color vision and had normal or corrected-to-normal visual acuity. Stimuli were displayed on a 21″ CRT monitor with a refresh rate of 140 Hz. Subjects sat with their head supported by a chin rest and viewed the monitor at a distance of 60 cm. Eye position was monitored online at 1000 Hz using an infrared eye tracker (SR Research, Canada).

In all experiments, a trial began with the presentation of a central fixation cross (white, 0.75° of visual angle) against a gray background. Once a stable fixation was recorded on the cross, a memory array was presented, consisting of a number of colored oriented bars (0.3° × 2°) randomly distributed around fixation at eccentricities in the range of 3°–6°, with a minimum center-to-center separation of 3° between items (example in Figure 1a). Each bar had a different color, and each bar's orientation was independently chosen at random from the full range of possible orientations (0°–180°).

The duration of display of the memory array was varied between trials and experiments. At the end of this exposure period, the memory array was replaced by a pattern mask, presented for 100 ms. A single (*probe*) bar was subsequently presented at fixation, with a color corresponding to one of the items in the preceding memory array. Subjects used an input dial (PowerMate USB Multimedia Controller, Griffin Technology, USA) to adjust the orientation of the probe item to match the remembered orientation of the item of the same color in the memory array (the *target*). The probe's initial orientation was randomly assigned. Accuracy was stressed, and responses were not timed. Any trial on which gaze deviated more than 2° from the central cross during presentation of the memory array was aborted and restarted with new feature values.

Analysis

A measure of recall *error* was obtained on each trial in each experiment by calculating the angular deviation between the orientation reported by the subject and the correct (target) orientation. For each combination of subject and display time (and cue validity in Experiments 2–4), we calculated the recall *bias*, defined as the mean of the recall error, and *precision*, defined as the reciprocal of the standard deviation of the error. As in previous studies (Bays et al., 2009; Bays, Wu, & Husain, 2011; Gorgoraptis, Catalao, Bays, & Husain, 2011), we used the definition of standard deviation for circular data given by Fisher (1995) and subtracted from the precision estimate the value expected by chance (i.e., if the subject had responded at random on each trial).

To quantify the contribution of different *sources of error* to overall precision estimates in each experiment, we applied a probabilistic model introduced by Bays et al. (2009; see also Zhang & Luck, 2008). This model attributes the distribution of responses on the reproduction
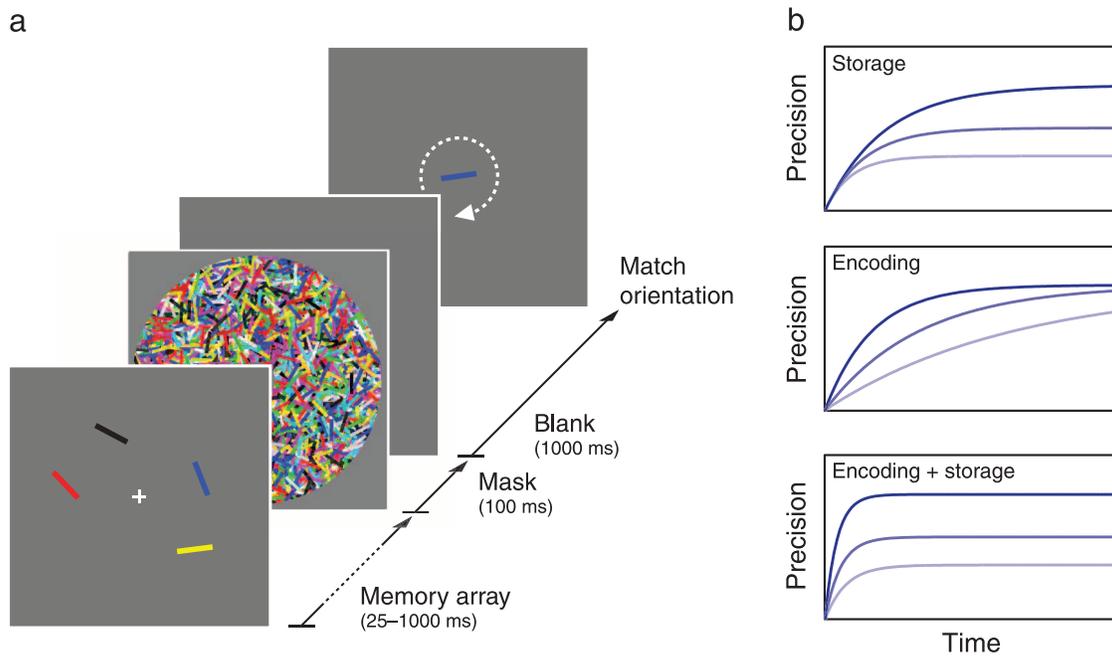
a

b

Figure 1. Assessing effects of exposure duration on recall of orientation. (a) The recall task used in Experiment 1. An array of colored oriented bars was presented for a variable exposure duration, followed by a pattern mask. After a blank retention interval, a *probe* bar appeared and subjects used a response dial to adjust its orientation to match the item with the same color in the memory array (the *target*). The angular difference between response and target orientations was taken as a measure of recall error. (b) Three hypotheses regarding the evolution of recall precision with exposure time, as a function of the number of items in the memory array. Here, lighter shades indicate more items stored in memory. The top panel is based on an assumption of limited *storage* capacity: as more items are stored, the maximum attainable precision declines. The middle panel depicts the case of limited *encoding* capacity: Eventually, the same level of precision is reached regardless of array size. Finally, the lower panel shows expected performance when both *storage and encoding* capacities are limited. Compare these possible results with the actual findings in Figure 2a.

task to a mixture of three components (illustrated in Figure 3a) corresponding to: reporting the target orientation (top), mistakenly reporting one of the other (non-target) orientations in the memory array (middle), and responding at random (bottom). Orientations of all memory array items are recalled with Gaussian variability.

Mathematically, the model is described by the following equation:

$$p\left(\hat{\theta}\right) = \alpha\phi_\kappa\left(\hat{\theta} - \theta\right) + \beta\frac{1}{m}\sum_i^m \phi_\kappa\left(\hat{\theta} - \varphi_i\right) + \gamma\frac{1}{2\pi}, \quad (1)$$

where θ is the true orientation of the target item, $\hat{\theta}$ is the orientation reported by the subject, and ϕ_κ is the von Mises distribution (the circular analogue of the Gaussian) with mean zero and concentration parameter κ. The probability of reporting the correct target item is given by α. The probability of mistakenly reporting a non-target item is given by β, and $\{\varphi_1, \varphi_2, \ldots \varphi_m\}$ are the orientations of the m non-target items. The probability of responding randomly is given by $\gamma = 1 - \alpha - \beta$.

Maximum likelihood estimates of the parameters α, β, γ, and κ were obtained separately for each subject and experimental condition using an expectation–maximization algorithm. The optimization procedure was repeated from a range of different initial parameter values to ensure that global maxima were obtained. Concentration κ was converted to the more familiar standard deviation, σ, according to the method of Fisher (1995).

Hypotheses regarding the effects of experimental parameters (exposure duration, array size) on recall precision and on each component of the mixture model were tested by analysis of variance (ANOVA). In Experiments 2–4, *t*-tests were used to test for precision advantages for valid over invalid or neutral trials and for neutral over invalid trials.

Analysis code is available online at http://www.sobell.ion.ucl.ac.uk/pbays/code/JV10/.

Experiment 1

This experiment investigated the observers' ability to reproduce from memory the orientations of objects presented in masked displays of varying size and duration. Previous studies testing recall of memory arrays with

unmasked displays have shown that the precision with which each visual item is stored declines rapidly as the number of items increases (Bays et al., 2009; Bays & Husain, 2008; Wilken & Ma, 2004). Figure 1b illustrates three hypotheses regarding the encoding of information into memory that are consistent with this finding.

One possibility (illustrated in the top panel) is that all items are initially encoded at a fixed rate, but a limit on memory capacity means that precision reaches a plateau at a maximum value that depends on the total number of items in the array. Thus, when two items are stored, each is recalled with lower precision than if only one is stored. An alternative possibility (middle) is that the maximum precision of storage is independent of memory load, but encoding is faster when there are fewer items; hence, precision may still depend on array size unless the display time is very long. The final possibility (bottom) is that both the maximum storage precision and the rate of encoding into memory depend on array size, with the two processes independently influencing precision of recall for any given display time.

Methods

Thirty-two subjects participated in Experiment 1. The procedure is illustrated in Figure 1a. The number of items in the memory array varied between subjects, with array sizes of 1, 2, 4, or 6 items, each tested in a different set of 8 subjects. Each item's color was chosen randomly from a set of easily distinguishable colors (black, white, red, green, blue, yellow, magenta, cyan). The duration of display of the memory array varied between trials. Each subject completed 800 trials in total, consisting of 100 randomly interleaved trials at each of 8 different display durations: 25, 50, 75, 100, 125, 300, 500, and 1000 ms. At the end of this exposure period, the memory array was replaced by a pattern mask for 100 ms, then by a blank interval of 1000 ms, followed by a single probe bar at fixation.

We calculated recall *bias* and *precision* for each combination of subject and display time based on error in reporting the target orientation. To capture the relationship between recall precision (P) and display duration (t), we fit an equation of the form:

$$P(t) = A(1 - e^{-t/\tau}). \tag{2}$$

A system obeying this RC equation (named after the Resistor–Capacitor electronics circuit that displays the same behavior) has the property that the *rate of increase* ($\dot{P}$) at time t is directly proportional to the difference between the current and maximum values (i.e., the "unfilled" capacity):

$$\dot{P}(t) = \dot{P}_0 \left(\frac{P_{max} - P(t)}{P_{max}} \right). \tag{3}$$

The temporal evolution of recall precision can, therefore, be fully described by just two parameters: the *maximum precision* P_{max}, and the *initial rate* $\dot{P}_0$. In previous studies (Bays et al., 2009; Bays & Husain, 2008) using unmasked displays, the relationship between precision of storage and the number of items stored was accurately captured by a power law: $P \propto N^{-\lambda}$. For each subject in the present study, we obtained least squares fits of a power law relationship between number of items (N) and each of the parameters describing precision (P_{max} and $\dot{P}_0$).

To examine the contribution different sources of error made to recall in this experiment, we fit a probabilistic model to the response data (see General methods section). Because the model fitting is data-intensive, and trials were divided between a large number of different memory array durations, for the purposes of this analysis, we binned trials into one of three *duration ranges*: short (25–50 ms), intermediate (75–300 ms), and long (500–1000 ms).

Results and discussion

The fidelity of recall can be characterized by two parameters, *bias* and *precision*. Bias indicates a systematic tendency to deviate from the correct target orientation in the same direction from trial to trial. No significant biases were observed for any array size or exposure duration ($p > 0.05$). Precision measures the degree to which responses cluster around the correct orientation: A precision of zero indicates that responses are randomly distributed relative to the target.

Consistent with previous findings, recall precision declined as the number of items in the memory array increased (Figure 2a, symbols; $F_{1,30} = 40.2$, $p < 0.001$). In addition, precision varied substantially with exposure duration ($F_{1,30} = 90.7$, $p < 0.001$). Examining performance for each set size independently (different colors in Figure 2a), a consistent pattern was observed, consisting of an initial rapid rise in recall precision as exposure was increased from the minimum duration, followed by saturation at longer display times as precision approached an asymptotic level.

More specifically, the relationship between precision and display time at each set size was accurately captured by an RC curve (Figure 2a, dotted lines), in which the rate of increase in precision at each time point is proportional to the remaining "unfilled" capacity, i.e., the difference between current and asymptotic precision values (see Methods section).

Each RC curve is described by two parameters: an *initial rate* (plotted in Figure 2b), corresponding to the slope of the curve at time zero, and a *storage limit* (Figure 2c), corresponding to the maximum value to which the curve asymptotes at long durations. Both the storage limit and the initial rate declined significantly with increasing number of items (limit: $F_{1,30} = 28.4$, $p < 0.001$; rate:
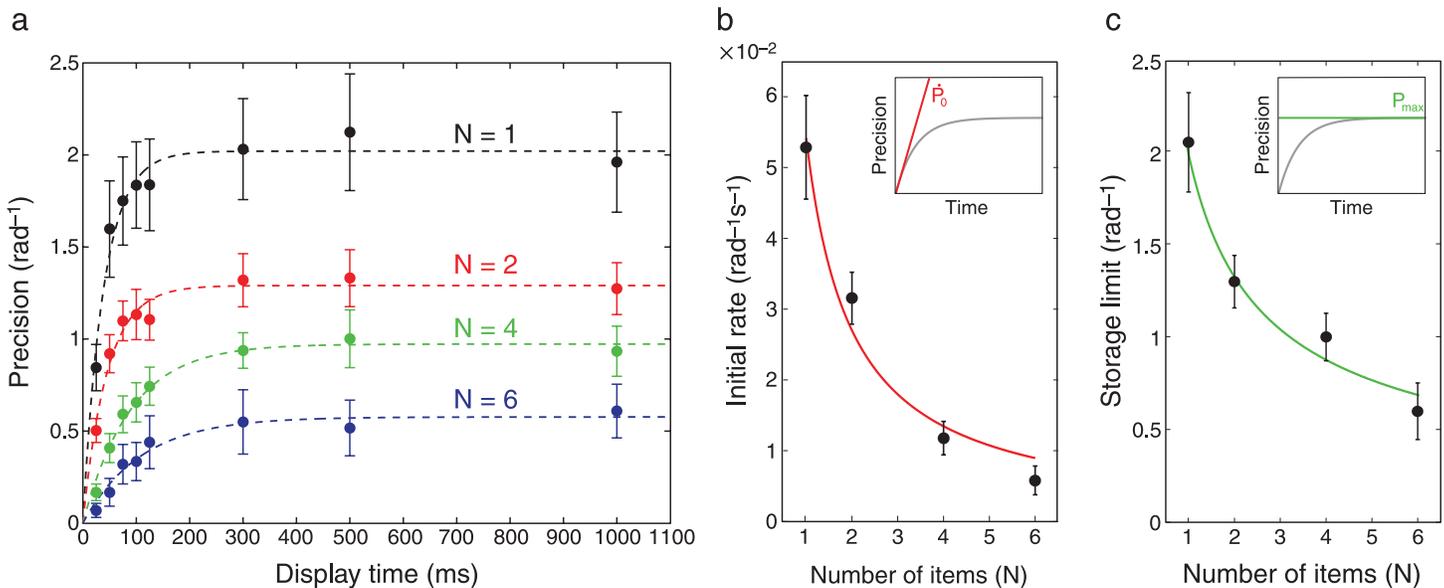
Figure 2. Temporal evolution of working memory precision. (a) Recall precision as a function of exposure duration and number of array items (*N*). Precision is defined as the reciprocal of the standard deviation of error in the reproduction task. Error bars indicate ±1 *SE*. Dashed lines indicate RC curves (see Methods section) with parameters that best fit the temporal evolution of recall precision at each array size. (b) The initial rate of encoding into memory was estimated from the rate parameter of the fitted RC curve (inset). The initial rate (black symbols) declined with array size according to a simple *inverse* relationship (i.e., rate proportional to $1/N$, red curve). (c) The upper limit on precision was estimated from the capacity parameter of each RC curve (inset). This storage limit (black symbols) declined with array size more slowly than encoding rate, following a *power law* relationship (i.e., maximum precision proportional to $N^{-\lambda}$, green curve).

$F_{1,30} = 52.4$, $p < 0.001$), consistent with the hypothesis illustrated in Figure 1b (bottom).

Figure 2b shows the relationship between *encoding rate* and the number of items in memory (*N*). A power law provided a good fit to the data ($R^2 = 0.93$), with the power constant being almost exactly one ($\lambda = 1.01 \pm 0.14$). Thus, this result is consistent with a simple inverse relationship between *encoding rate* (initial rate of change of precision, $\dot{P}_0$) and array size, i.e., $\dot{P}_0 \propto 1/N$.

By contrast, the relationship between number of items and maximum precision (the *storage limit*) was very different. Previous studies (Bays et al., 2009; Bays & Husain, 2008) using long exposures have observed a *power law* relationship ($P \propto N^{-\lambda}$) relating the number of items in an array (*N*) to the precision with which each item is stored (*P*). A similar relationship was observed in the present study (green line, Figure 2c; $R^2 = 0.96$). However, although the power constant that best fit the data ($\lambda = 0.60 \pm 0.12$) was consistent with that obtained in a previous study based on an orientation discrimination task ($\lambda = 0.69$; Bays & Husain, 2008), it was significantly different from that between number of items and rate of encoding ($t_6 = 3.7$, $p = 0.007$).

Thus, although both storage and encoding into working memory can be related to *N* by a power function, encoding rate fell more rapidly with increasing number of items (compare Figures 2b and 2c).

The precision measure used thus far to describe performance is a non-parametric statistic reflecting the fidelity of recall of the target orientation, independent of any particular model of the underlying response distribution. To investigate how different possible sources of error contribute to memory precision at different exposure durations, we applied to the data a previously developed probabilistic model (Bays et al., 2009, 2011; Gorgoraptis et al., 2011) that describes the response distribution in terms of three different types of error (illustrated in Figure 3a; see General methods section for details).

The first source of error is Gaussian variability in recall of the target orientation (Figure 3a, top). Black symbols in Figure 3b plot the standard deviation of the Gaussian error component for each set size as a function of display duration (*short*: ≤50 ms; *intermediate*: 75–300 ms; *long*: ≥500 ms). Consistent with previous results, recall variability increased significantly with increasing number of items ($F_{1,30} = 5.8$, $p = 0.022$). Furthermore, recall variability *decreased* significantly with increasing exposure time ($F_{1,30} = 5.0$, $p = 0.034$).

A second source of error arises in multiple-item arrays, where subjects on occasion mistakenly report the orientation of one of the non-target items in the preceding array (Figure 3a, middle). Results of a previous study (Bays et al., 2011) indicated that these errors are due to *misbinding*, i.e., errors in recalling which color belongs
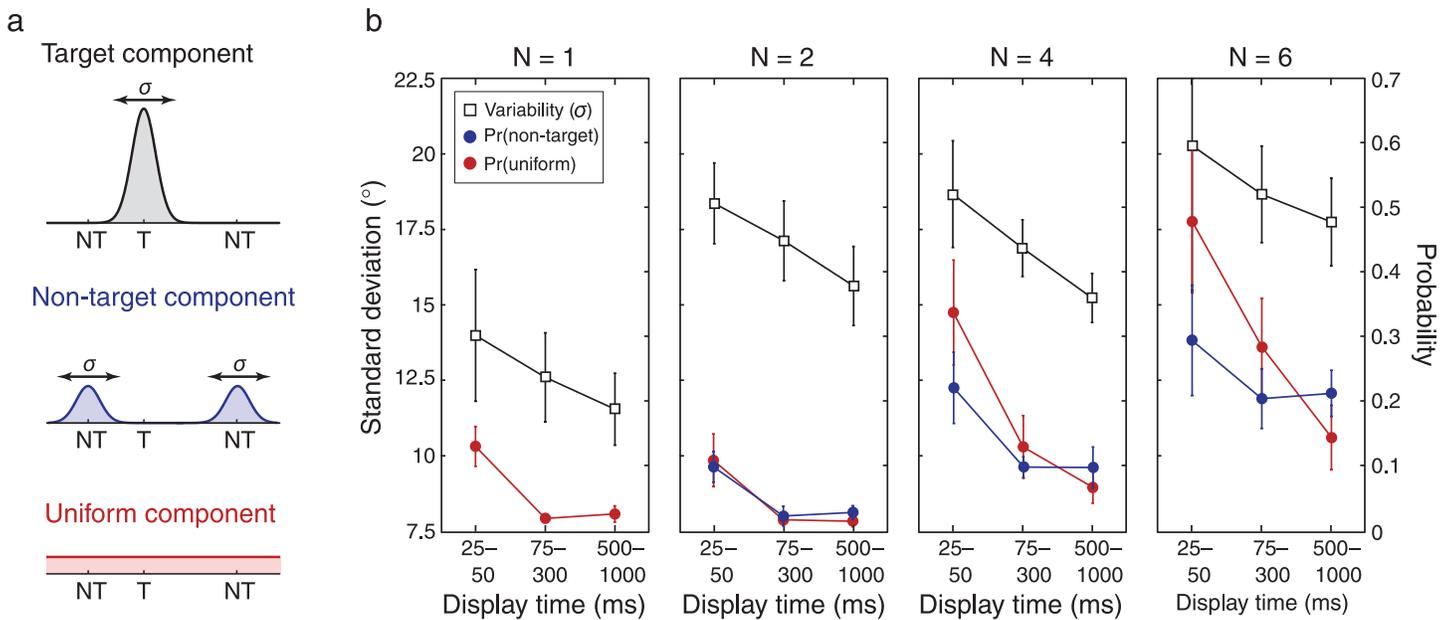
Figure 3. Components of error in the working memory task. (a) The distribution of responses was decomposed into a mixture of three separate components: responses distributed with Gaussian variability around the correct (target, T) orientation (top), responses distributed around the orientations of other, non-target (NT) items in the memory array (middle), and random responses distributed uniformly throughout the response space (bottom). (b) Maximum likelihood estimates of parameters of the mixture model illustrated in (a), for different array sizes and exposure durations. Gaussian variability in target responses (black symbols) increases with array size but decreases with exposure duration. The frequencies of non-target (blue) and random (red) responses also increase with array size. Note that random responding declines rapidly with increasing exposure duration, but non-target (*misbinding*) errors maintain a constant frequency at intermediate and long durations. Error bars indicate ±1 *SE*.

with which orientation (Allen, Baddeley, & Hitch, 2006; Robertson, 2003; Treisman, 1998; Wheeler & Treisman, 2002; Wolfe & Cave, 1999). The frequency of binding errors is shown by the blue symbols in Figure 3b. Consistent with previous results, misbinding frequency increased with increasing memory load ($F_{1,30} = 49.2$, $p < 0.001$). Binding errors were also more frequent at the briefest display times (*short* vs. *intermediate*: $F_{1,30} = 8.7$, $p = 0.006$), but their frequency subsequently appeared to plateau at a value determined by array size (*intermediate* vs. *long*: $F_{1,30} = 0.7$, $p = 0.73$).

A final source of error corresponds to responses randomly distributed relative to target and non-target orientations (Figure 3a, bottom). These errors may occur when no information about the target orientation has been stored and subjects simply "guess" at random. Like misbinding errors, random responses (red symbols in Figure 3b) increased in frequency with increasing number of items in the array ($F_{1,30} = 21.1$, $p < 0.001$), making a substantial contribution to the response distribution at the largest set sizes (48% of responses for 6-item arrays at the shortest display times).

Crucially, however, the frequency of random responses fell rapidly toward zero as exposure duration increased, with no indication of a plateau as observed for misbinding (*short* vs. *intermediate*: $F_{1,30} = 36.8$, $p < 0.001$;

intermediate vs. *long*: $F_{1,30} = 9.4$, $p = 0.005$). This finding is important for comparison with studies that have attempted to measure WM storage capacity based on very briefly displayed (≤200 ms) memory arrays (e.g., Anderson, Vogel, & Awh, 2011; Zhang & Luck, 2008; see General discussion section).

In this experiment, we used presentation of a masking stimulus as a probe into the time course of WM encoding. One assumption of this approach is that replacing the memory array with a pattern mask halts encoding of the memory items but does not significantly disrupt the visual information that has already entered working memory storage. An alternative possibility is that there exists a "window of integration" over which period visual information from the array and the mask cannot be fully differentiated, resulting in a noisy representation of the array entering memory.

This possibility was convincingly addressed by a previous study using the same masking technique to probe WM (Vogel et al., 2006). If the effects of masking a briefly presented array are due to sensory integration, it should not matter in which order the mask and array are presented; however, these authors showed that a mask that strongly disrupted recall when presented *after* the array had no effect on performance when presented *before*. This strongly supports the conclusion that performance costs

associated with presenting a mask after the array (as in the present study) are primarily due to halting encoding before WM capacity is reached rather than sensory integration between the mask and array.

Experiment 2

Previous studies have demonstrated that memory resources can be flexibly allocated to prioritize storage of a cued array item (Bays & Husain, 2008; Gorgoraptis et al., 2011). Experiment 2 was designed to investigate the time course of this cuing effect.

Methods

Sixteen subjects participated in Experiments 2A and 2B (8 in each). Each trial began with presentation of a fixation cross, followed by a memory array, as in Experiment 1. The memory array on each trial consisted of two colored bars (one blue, one red) with randomly selected locations and orientations. A white disk (the *cue*) was briefly flashed at the location of one of the two items simultaneously with the onset of the memory array (25-ms duration, 2.5° diameter; example in Figure 4a).

The duration of the memory array varied between trials, with each subject completing 150 trials at each of 4 different display durations (100, 200, 400, 800 ms), randomly interleaved. This allowed us to examine the consequences of initial cuing on performance over time. At the end of this exposure period, the memory array was replaced by a pattern mask for 100 ms, followed by a single probe bar at fixation. Subjects adjusted the orientation of the probe bar to match the item of the same color in the memory array. Central fixation was enforced as above.

For subjects in Experiment 2A, the cue was *predictive* of the subsequent probe: On two out of three trials (randomly interleaved), the color of the probe item corresponded to the memory item cued by the white disk. For subjects in Experiment 2B, the cue was *non-predictive*: The probe item was equally likely to have the color of either item in the memory array.

Note that, in line with previous studies (e.g., Vogel et al., 2006), Experiment 1 included a blank retention interval following the mask, before presentation of the probe stimulus. In tasks examining recall of unmasked stimuli, such a retention interval is necessary to prevent short-lived iconic memory traces contributing to performance,
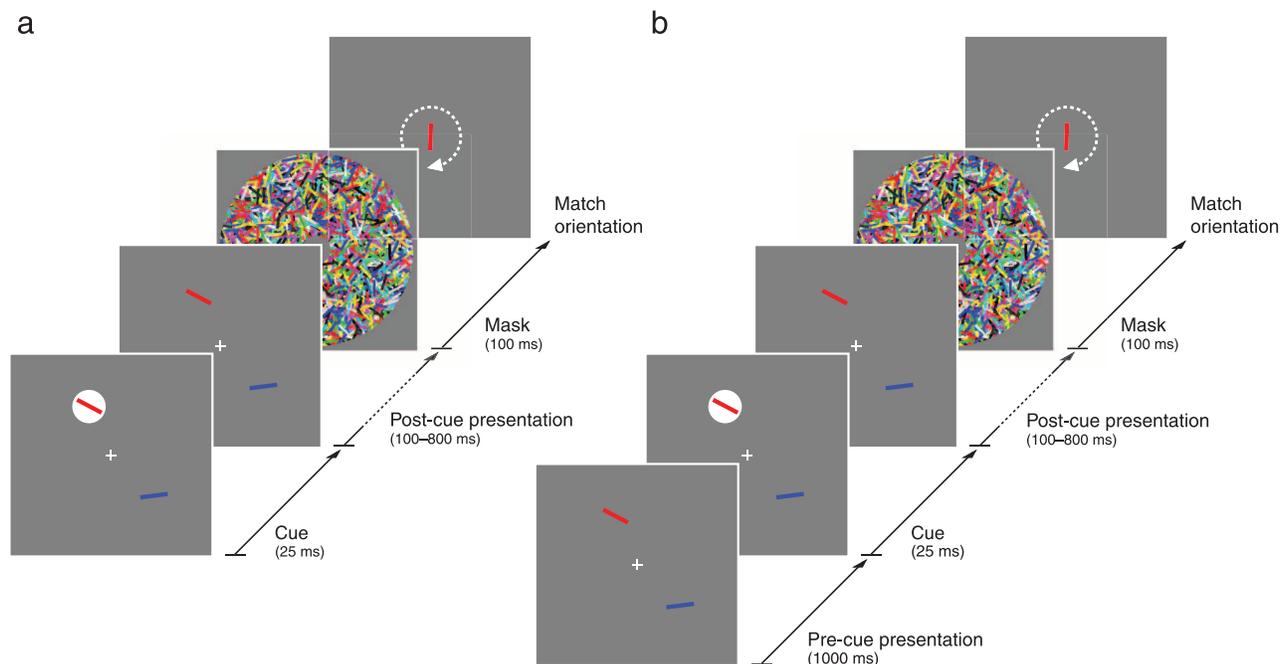

Figure 4. Assessing cuing effects on recall performance. (a) The recall task used in Experiment 2. Memory arrays consisted of two randomly oriented bars (one red, one blue). A randomly selected bar was cued by a briefly flashed white disk presented simultaneously with the onset of the memory array. After a variable post-cue display period, in which the memory array remained visible, a pattern mask was presented. Subjects adjusted a probe bar to reproduce the orientation of one of the array items, as in Experiment 1. (b) The recall task used in Experiment 3. The procedure was identical to Experiment 2, except that the memory array was displayed for 1000 ms *before* one of the bars was cued, in addition to the variable post-cue display period. In Experiments 2A and 3A, the cue predicted which item would be probed on 2/3 trials; in Experiments 2B and 3B, the cue was not predictive of the probe.

but as iconic memory is also erased by a pattern mask (Coltheart, Laming, Routh, & Broadbent, 1983; Sperling, 1960), this delay is not essential to the present experiments. Therefore, to ensure any cuing effects observed in Experiment 2 reflected allocation of resources during array presentation rather than processes occurring during maintenance (e.g., shifts of attention within the memory representation; Griffin & Nobre, 2003), the retention interval was not included in Experiment 2 or the subsequent cuing experiments.

Results and discussion

In this experiment, subjects were tested on their recall of two-item memory arrays, similar to those in Experiment 1.
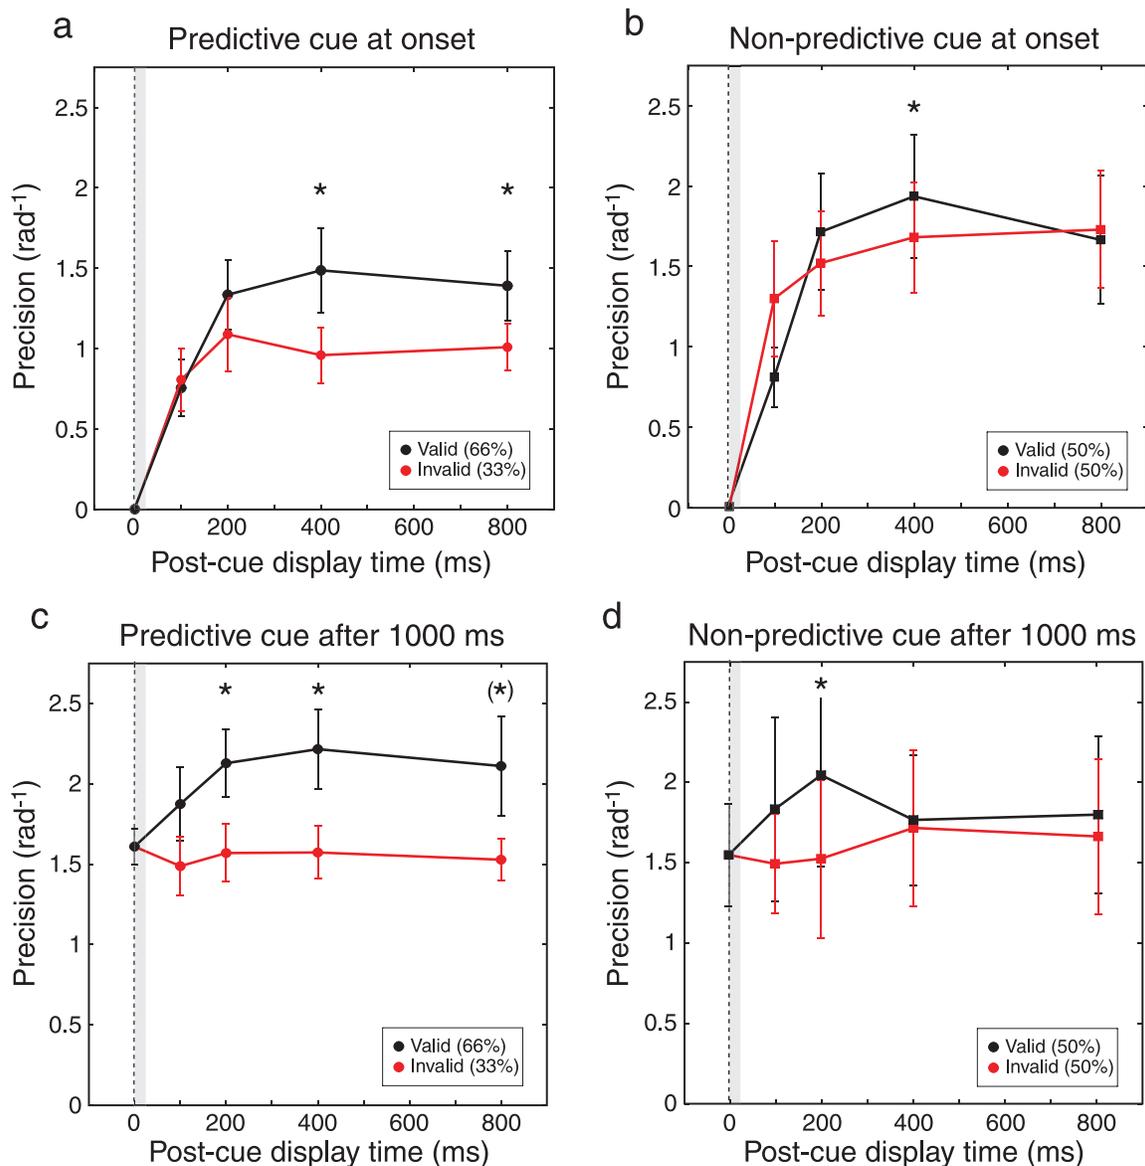

Figure 5. Time course of cuing effects on encoding and maintenance. (a) Recall precision as a function of exposure duration when a *predictive* cue was presented at array onset, for trials on which memory was probed for the cued item (*valid* trials, black) and non-cued item (*invalid* trials, red). Error bars indicate ± 1 SE. Asterisks indicate time points at which there was a significant recall advantage for cued over non-cued items ($p < 0.05$). Dotted vertical line and shaded area indicate onset and duration of the cue event. (b) Recall precision plotted as in (a) but for subjects presented with non-predictive (task-irrelevant) cues. Note that, unlike for predictive cuing, the advantage for the cued item is abolished at longer exposure durations. (c) Recall precision for *predictive* cues presented after 1000-ms exposure to the memory array, plotted as a function of *post-cue* display time. Note that high-resolution representations of both items are already stored in memory at the time of cue presentation (0-ms condition). Asterisk indicates a cued item advantage with borderline statistical significance ($p = 0.054$). (d) Recall precision for *non-predictive* cues after 1000-ms exposure. Note that the cued item advantage is again abolished at longer exposures.

Simultaneous with the onset of the memory array, one item was highlighted by a very brief flash (the *cue*; Figure 4). The effect of this cuing event on the subsequent encoding and storage of the array items was examined by comparing recall performance for the cued and non-cued items (Figure 5). The post-cue display time was varied to assess the evolution of cued/non-cued differences in the time following the cue event.

Figure 5a shows recall precision for cued (black) and non-cued items (red) in Experiment 2A, in which the cue was *predictive* of which item would be probed (cued items were probed twice as frequently as non-cued items). No differences in recall precision were observed when exposure time was brief (≤ 200 ms; $t_7 < 1.1$, $p > 0.16$), but for display durations of 400 ms and longer, a significant recall advantage was observed for the cued item ($t_7 > 2.3$, $p < 0.03$).

Figure 5b shows performance in Experiment 2B, which was identical except that the cue was *non-predictive* of the probe (recall was tested for cued and non-cued items with equal frequency). In this case, a small recall advantage for the cued item was observed at 400 ms ($t_7 = 2.2$, $p = 0.034$), but the effect was short-lived, with no evidence for a difference in recall at 800 ms ($t_7 = 0.5$, $p = 0.61$).

The observation of a precision advantage for the flashed item even when the flash was irrelevant to the memory task implies that a recall advantage can develop "bottom-up" in response to salient elements in the visual environment. However, the effect of task-irrelevant cues was rapidly abolished with continuing presentation of the array, while the effect of predictive cues was maintained, consistent with an additional "top-down" mechanism enhancing storage of items that are relevant to current goals.

Experiment 3

Because the cue was presented simultaneously with onset of the memory array in Experiment 2, the observed effects on recall could reflect either preferential encoding or preferential allocation of memory to the cued item. To distinguish between these possibilities, in Experiment 3, the cue was presented only after the memory array had been visible for 1000 ms (Figure 4b), sufficient time for both items to be fully encoded into memory.

Methods

Twelve subjects participated in Experiments 3A and 3B (6 in each). The protocol was identical to Experiment 2 except that the memory array was displayed for 1000 ms *before* cue onset, so encoding of the array itself could not be a limiting factor on performance (Figure 4). The cue was displayed for 25 ms, and subsequent duration of the memory array *following* the cue event varied between trials as above. Each subject completed 200 trials at each of 4 different post-cue display durations (100, 200, 400, 800 ms) as well as 200 trials in a baseline (0 ms) condition where the mask was displayed at 1000 ms without a preceding cue event. All trial conditions were randomly interleaved. As above, the cue was predictive of the subsequent probe in Experiment 3A and non-predictive in Experiment 3B.

Results and discussion

Figure 5c shows recall precision for cued and non-cued items when a *predictive* cue was presented after 1000 ms as a function of the subsequent *post-cue* display time.

As expected, in a baseline (0 ms) condition with no cue, array items were recalled with similar precision (1.6 rad^{-1}) to that observed in Experiment 1 for the same array size and display time (1.3 rad^{-1}). This is consistent with our prediction that encoding of both items into memory would be completed by the time of cue presentation in this experiment. Nonetheless, following presentation of the cue, a significant recall advantage developed for the cued item (Figure 5c). Note that this effect of cuing was present at a shorter post-cue exposure duration than when the cue was presented at onset (200 ms: $t_5 = 2.6$, $p = 0.02$), as well as being observed at subsequent time points (400 ms: $t_5 = 3.0$, $p = 0.015$; 800 ms: $t_5 = 1.9$, $p = 0.054$).

The effects of a *non-predictive* cue at 1000 ms are shown in Figure 5d. A significant cue advantage was observed, again at a shorter exposure duration (200 ms: $t_5 = 3.3$; $p = 0.023$) than for a cue at onset. However, as when a non-predictive cue was presented at onset (Figure 5b), the advantage for the cued item was abolished at longer post-cue display times (400 ms: $t_5 = 0.3$, $p = 0.40$; 800 ms: $t_5 = 1.3$, $p = 0.13$).

The differences in time course between predictive and non-predictive cue effects confirm the findings of Experiment 2 that recall precision is influenced by both "bottom-up" (salience-based) and "top-down" (goal-based) mechanisms. The present results further demonstrate that the enhancement of recall for cued items can occur even once both items have been fully encoded into memory, indicating a *reallocation* of previously allocated memory resources to reflect the change in relative priority of the stimuli.

To investigate which of the possible sources of error was responsible for the difference in recall precision between cued and non-cued items, we applied the probabilistic model illustrated in Figure 3a to the data from Experiments 2 and 3. A significant interaction between cue validity and post-cue duration was observed for the standard deviation of the Gaussian error component ($F_{1,19} = 6.7$, $p = 0.018$), indicating that variability was lower for cued than non-cued items at longer post-cue durations. No significant validity or validity $\times$ duration

effects were observed for non-target ($F_{1,19} = 2.4$, $p = 0.14$) or uniform ($F_{1,19} = 0.6$, $p = 0.45$) components of the model. Hence, the effects of cuing do not reflect changes in the probability that the cued item is stored but rather the resolution of its storage.

Experiment 4

The results of Experiments 2 and 3 suggest that working memory resources can be rapidly *reallocated* to a cued item in order to store it with greater fidelity. A strong prediction of this hypothesis is that the increase in recall precision for a cued item will coincide with a *decrease* in precision for other items, which now receive a smaller proportion of memory resources. In order to test this hypothesis, in Experiment 4, we interleaved *cue* trials, in which we predicted that memory resources would be preferentially allocated to the item highlighted by the flash, with baseline (*neutral cue*) trials designed to encourage equal distribution of resources between items.

Methods

Eight subjects participated in Experiment 4. Experimental parameters were chosen to maximize cuing effects, based on results of the previous experiments: Cues were *predictive* of the probe and presented after 1000 ms, followed by a post-cue exposure duration of 400 ms. The protocol was identical to Experiment 3A,

except for the addition of a *neutral cue* condition in which white disks were flashed at the locations of *both* array items. Both items were flashed on these trials to control for general alerting effects of the cue event, independent of which item is cued. Each subject completed 608 trials in total, consisting of 304 *valid cue* trials (probe matched the cued item), 152 *invalid cue* trials (probe matched the non-cued item), and 152 *neutral cue* trials (both items were cued), randomly interleaved.

Results and discussion

As expected, cuing condition had a substantial effect on the precision with which array items were recalled (Figure 6a). Recall for cued items, on *valid* trials, was significantly better than the baseline precision measured on *neutral cue* trials ($t_7 = 2.1$, $p = 0.04$). Consistent with the resource reallocation hypothesis, baseline performance was, in turn, superior to recall of non-cued items, measured on *invalid* trials ($t_7 = 3.7$, $p = 0.004$).

Figure 6b shows the best-fitting parameters of the probabilistic model (Figure 3a) obtained for responses in each cuing condition. Non-target (blue symbols) and random responses (red symbols) made negligible contribution to the response distribution (<3% of responses) and did not vary between cue conditions ($F_{2,14} < 2.4$, $p > 0.13$). In contrast, variability in the Gaussian response component varied significantly with cue condition ($F_{2,14} = 5.6$, $p = 0.016$), in a manner consistent with the overall effect on precision (*valid* < *neutral*: $t_7 = 2.1$, $p = 0.04$; *neutral* < *invalid*: $t_7 = 1.8$, $p = 0.05$).
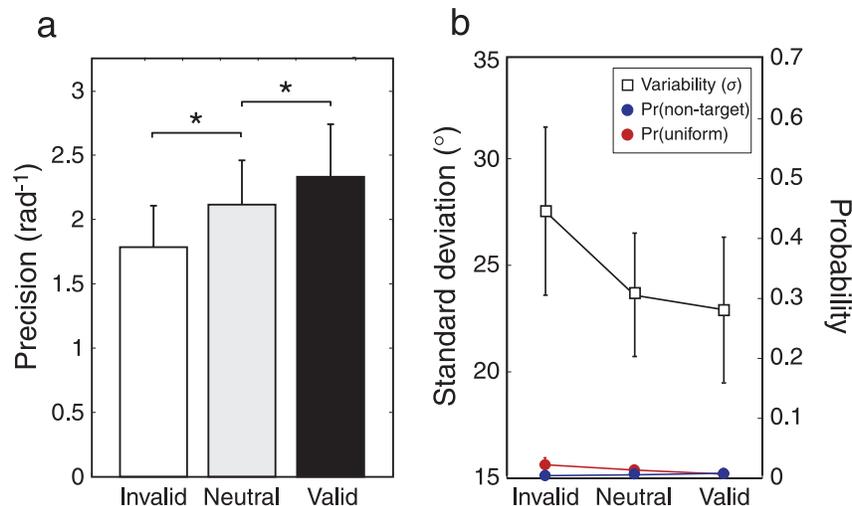

Figure 6. Recall of cued and non-cued items relative to baseline. (a) Recall precision in Experiment 4 on *invalid* and *valid cue* trials, in which one array item was cued with a flash, and on *neutral cue* trials, in which *both* items were flashed. Cues were presented after 1000 ms, and post-cue exposure was 400 ms. Single cues were predictive of which item would be probed (*valid* trials twice as frequent as *invalid*). Asterisks indicate significant differences, $p < 0.05$. (b) Maximum likelihood estimates of parameters of the mixture model (illustrated in Figure 3), for *invalid*, *valid*, and *neutral cue* trials. Note that cuing condition is reflected in variability of the Gaussian response component, while non-target and uniform components make negligible contribution to errors.

These results confirm that the increase in storage resolution observed for the cued item comes at the cost of a decrease in resolution for the uncued item, consistent with a partial withdrawal of limited memory resources from the low-priority item and their reallocation to the item with greater relevance to the task.

General discussion

Previously, studies of working memory encoding have been based on techniques such as change detection (Luck & Vogel, 1997; Pashler, 1988; Vogel et al., 2006; Woodman & Vogel, 2005) or whole report (Gegenfurtner & Sperling, 1993; Shibuya & Bundesen, 1988; Sperling, 1960), which are intended to measure *whether or not* a stimulus is stored in WM. However, a growing body of evidence indicates that this binary classification is insufficient as a description of WM storage, because stored items can also vary in the *fidelity* of their representation (Alvarez & Cavanagh, 2004; Awh, Barton, & Vogel, 2007; Bays et al., 2009; Bays & Husain, 2008; Bays et al., 2011; Fougnie et al., 2010; Gorgoraptis et al., 2011; Lakha & Wright, 2004; Palmer, 1990; Wilken & Ma, 2004; Zhang & Luck, 2008). In the present study, we combined masked presentation of a memory array with a reproduction task in order to assess the *precision* with which briefly presented visual information is represented in WM.

Time course of encoding and storage in working memory

In Experiment 1, we observed effects of presentation duration on recall precision (Figure 2) that were accurately described by the interaction of two separate constraints: a processing limit that determines the *encoding rate* at which visual input enters WM and a *storage limit* that sets the maximum precision with which this input can be maintained.

Following prolonged exposure to a visual array, recall should be constrained by limits on storage only. At the longest display durations, the precision with which each visual item was recalled declined monotonically with the number of items in the array. As observed in previous studies (Bays et al., 2009; Bays & Husain, 2008), this relationship between precision and memory load was accurately described by a power law (Figure 2c). The results are consistent with a limit on the total amount of visual information that can be maintained in WM.

In contrast, when a visual array is presented very briefly, the quality of subsequent recall should depend primarily on how rapidly visual information can be transferred into WM. At the shortest exposure durations, the precision of recall was significantly reduced compared

to prolonged exposure, but precision was still highly dependent on the total number of items to be encoded. Rather than a power law, the initial rate of rise in precision was found to have a simple inverse ($1/N$) relationship with the number of items in the array (Figure 2b), consistent with a limit on processing capacity that is independent of the previously identified limit on storage.

This same inverse relationship between processing rate and set size is the basis of an influential model of visual attention (TVA), which has had considerable success in reproducing many classical results in visual search and divided attention (Bundesen, 1990; Bundesen & Habekost, 2008). TVA is an example of a *parallel model*, in which multiple stimuli are processed simultaneously in a race for storage in working memory, as distinct from *serial models* in which stimuli are selected one at a time for attentional processing followed by transfer to WM (Hoffman, 1979; Wolfe, 1994).

Sources of error in recall: Variability, misbinding, and guessing

Both parallel and serial models are consistent with an increase in recall performance with exposure duration, as observed in the present study, but they make different predictions regarding the distribution of errors in the reproduction task. To investigate, we applied to our data a probabilistic model of response selection (Bays et al., 2009), which assigns errors to one of three components (Figure 3a): Gaussian-distributed errors due to variability in recall of the target orientation, "binding errors" where the orientation of one of the other non-target items is erroneously reported, and random errors that are unrelated to any of the orientations in the memory array.

According to a serial model of WM encoding, reducing the encoding time results in a step-wise decrease in the *number of items* present in memory and, hence, an increase in random responding (i.e., guessing). However, each item that gains access to memory has already been fully processed to the maximum possible resolution, so this model predicts no effect of exposure duration on the *variability* of the Gaussian error distribution.

In fact, both random responding *and* variability declined with increasing exposure duration (Figure 3b). The changes in variability rule out a strictly serial model, in which items are transferred to WM only once their processing is complete. Indeed, substantial effects of exposure duration on recall variability were observed even when there was only one item in the display, indicating gradual accumulation of visual information into a single WM representation.

At the briefest presentation times, there was evidence for a substantial frequency of random responses, particularly with larger arrays. While superficially this appears inconsistent with a process in which all array items are encoded simultaneously into WM, parallel models such as

TVA typically incorporate a stochastic element, such that different items are encoded at different rates from trial to trial (Bundesen, 1990). As a result, at any given moment during encoding, each item will be at a different stage of representation in working memory. For very brief presentations, therefore, encoding of some items may be at such an early stage that responses are indistinguishable from chance.

Consistent with this hypothesis, random responses declined to negligible levels as exposure time increased (Figure 3). At the longest exposures, the frequency of random responding was equivalent to less than one item per array, indicating that information eventually accumulated in WM about *every* item presented. However, the evidence for incomplete encoding of larger arrays even with exposures as long as 300 ms may have important consequences for studies that attempt to measure WM storage capacity based on very briefly displayed memory arrays.

Our findings suggest that recall errors following brief exposures may reflect incomplete encoding of array stimuli into working memory, as opposed to limits on its capacity. This may, in part, explain why some recent studies using exposures of ≤200 ms have obtained results that, interpreted in terms of capacity, would indicate that working memory can hold only ~2 colors or orientations at one time (Anderson et al., 2011; Zhang & Luck, 2008), a finding not supported by the present results nor by previous studies using >500 ms presentation times (e.g., Bays et al., 2009; Fougnie et al., 2010; Gorgoraptis et al., 2011).

The habitual use of short exposure durations in working memory studies appears to have arisen with two aims in mind: (1) minimizing eye movements during the exposure period and (2) limiting the contribution of long-term memory (LTM) to recall. Shifts of gaze between display elements could potentially disrupt encoding through saccadic suppression of visual input or bias storage toward the smaller set of fixated memory items. However, these effects are best excluded by active monitoring of eye position (as in the present study) rather than reducing display duration.

The possibility that activation of long-term memory representations may contribute to recall performance is more difficult to exclude. However, selecting memory targets from a continuous feature space (e.g., the 180° range of orientations in the present study) as opposed to a small set of discrete features values (as common in change detection tasks, e.g., Luck & Vogel, 1997) will tend to limit the usefulness of LTM representations as aides to recall. More importantly, there is little evidence to suggest that contamination by LTM is prevented by using brief (e.g., 100 ms) displays. In particular, the smooth evolution of recall precision with time observed in the present study (Figure 2a) does not support a two-stage process of sequential encoding into WM and then LTM but rather is consistent with a continuous process of encoding visual information into a single capacity-limited store.

The third component of the response distribution corresponds to binding errors (Allen et al., 2006; Bays et al., 2011; Robertson, 2003; Treisman, 1998; Wheeler & Treisman, 2002; Wolfe & Cave, 1999). These errors occur because accurate reproduction of the probed item requires not only recall of the orientations in the preceding array but also recall of which orientation corresponds to the probed color. If the "binding information" that pairs orientations with colors becomes corrupted, subjects may respond with the orientation of one of the other, non-target items in the memory array.

These binding errors occurred with greatest frequency at the shortest exposure durations, but importantly, unlike random errors, their frequency appeared to plateau at a constant level as presentation time increased. This limiting frequency increased with array size, indicating that binding errors became more prevalent as array size increased. Hence, the overall decline in recall performance with increasing memory load has two main components: an increase in the variability with which individual features are stored and a decline in the fidelity with which bindings between feature dimensions are maintained. Similarly, both misbinding and increasing variability contribute independently to the increase in recall error observed at shorter encoding times (Figure 3b).

The observation of increases in both variability and misbinding with number of items in the memory array, here and in previous studies (Bays et al., 2009; Bays & Husain, 2008, 2009; Bays et al., 2011; Fougnie et al., 2010; Gorgoraptis et al., 2011), has been interpreted as supporting a *shared-resource* model of visual working memory. According to this proposal, a single memory resource is distributed between the elements of a visual scene. As more items are stored, less resource is available per item, with the result that both object features and feature bindings are maintained with decreasing fidelity.

Time course of allocation and reallocation of working memory resources

A critical claim of the shared-resource account of WM is that the distribution of resources is *flexible*, such that prioritized items can be allocated a greater share of resources and, hence, be remembered with enhanced precision (Bays & Husain, 2008). This marks the clearest point of contrast with the influential "slot" model of WM, in which every visual object is either represented in its entirety in an individual memory slot or else not stored at all (Cowan, 2001; Luck & Vogel, 1997; Pashler, 1988).

In Experiments 2–4, we observed significant advantages in the precision of recall for a memory array item cued by a brief flash (Figures 4 and 5). This effect was due to a decrease in the variability with which the cued item was

stored rather than changes in the probability of storage (Zhang & Luck, 2008), and this decrease was matched by a corresponding *increase* in variability for the non-cued item (Figure 6), as predicted by a resource model.

A precision advantage for the flashed item was observed even when the flash was irrelevant to the memory task, indicating an automatic "bottom-up" allocation of memory resources to the location of the salient external event. However, varying the post-cue presentation time of the memory array revealed that this bottom-up effect was short-lived: Precision of cued and non-cued items returned to equality within a few hundred milliseconds given continued exposure to the array.

In contrast, if the flash was predictive of which item was most likely to be probed in the subsequent test of recall, the advantage for the cued item was maintained even up to the longest post-cue exposure durations (800 ms). This is consistent with the influence of an additional "top-down" influence on working memory allocation, biasing the resource distribution toward the goal-relevant item.

Crucially, while the importance of both bottom-up and top-down influences on *encoding* of visual information is widely recognized (Bays, Singh-Curry, Gorgoraptis, Driver, & Husain, 2010; Bundesen, 1990; Desimone & Duncan, 1995; Kastner & Ungerleider, 2000; Lamy & Zoaris, 2009; Posner, 1980; Theeuwes & Burger, 1998), the present effects on recall precision were observed even when the flash occurred during the *maintenance* phase of working memory storage, i.e., once both items had been fully encoded into memory (Figures 5c and 5d). Indeed, even when the cue was presented at the onset of the array, significant advantages for the cued item did not develop until after encoding of both items was largely complete (with the result that the cue effect was delayed compared to presentation during maintenance).

These results are consistent with a rapid *reallocation* of limited working memory resources. Storage capacity, initially equally distributed between items, is partially withdrawn from the uncued item, with a cost to the resolution of its representation in memory. This freed capacity is used to encode additional information about the cued item, enhancing the precision of its memory representation.

Conclusions

In this study, we have examined how the precision with which visual objects are stored in working memory depends on the duration of their presentation. Recall performance was limited by the *rate* at which visual information could be encoded into memory. Our findings are consistent with a parallel encoding model in which multiple items are processed simultaneously, resulting in increasingly precise representation in WM as exposure time increases.

Recall was also constrained by an upper limit on the information stored about the array, consistent with allocation of a *limited resource* of WM storage determining the maximum resolution with which each item can be maintained. Cuing individual items within the array revealed flexible reallocation of storage, increasing the resolution of recall for visually salient or behaviorally important items at the cost of reduced precision for lower priority items. Such redistribution may allow an optimal allocation of memory resources to be maintained in the face of frequent shifts in the behavioral relevance of objects in our visual environment.

References

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,* 298–313.

Alvarez, G. A., & Cavanagh, P. (2004). The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychological Science, 15,* 106–111.

Anderson, D. E., Vogel, E. K., & Awh, E. (2011). Precision in visual working memory reaches a stable plateau when individual item limits are exceeded. *Journal of Neuroscience, 31,* 1128–1138.

Awh, E., Barton, B., & Vogel, E. K. (2007). Visual working memory represents a fixed number of items regardless of complexity. *Psychological Science, 18,* 622–628.

Baddeley, A. D., & Hitch, G. J. (1974). Working memory. In K. W. Spence & J. T. Spence (Eds.), *The*

psychology of learning and motivation (pp. 47–89). New York: Academic Press.

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):7, 1–11, http://www.journalofvision.org/content/9/10/7, doi:10.1167/9.10.7. [PubMed] [Article]

Bays, P. M., & Husain, M. (2008). Dynamic shifts of limited working memory resources in human vision. *Science, 321,* 851–854.

Bays, P. M., & Husain, M. (2009). Response to comment on "Dynamic shifts of limited working memory resources in human vision". *Science, 323,* 877d.

Bays, P. M., Singh-Curry, V., Gorgoraptis, N., Driver, J., & Husain, M. (2010). Integration of goal- and stimulus-related visual signals revealed by damage to human parietal cortex. *Journal of Neuroscience, 30,* 5968.

Bays, P. M., Wu, E. Y., & Husain, M. (2011). Storage and binding of object features in visual working memory. *Neuropsychologia, 49,* 1622–1631.

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, 1–34, http://www.journalofvision.org/content/11/5/4, doi:10.1167/11.5.4. [PubMed] [Article]

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

Bundesen, C. (1990). A theory of visual attention. *Psychological Review, 97,* 523–547.

Bundesen, C., & Habekost, T. (2008). *Principles of visual attention: Linking mind and brain*. Oxford: Oxford University Press.

Chun, M. M., & Potter, M. C. (1995). A two-stage model for multiple target detection in rapid serial visual presentation. *Journal of Experimental Psychology: Human Perception and Performance, 21,* 109.

Coltheart, M., Laming, D. R. J., Routh, D. A., & Broadbent, D. E. (1983). Iconic memory. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 302,* 283–294.

Cowan, N. (1995). *Attention and memory: An integrated framework*. USA: Oxford University Press.

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences, 24,* 87–114.

Desimone, R., & Duncan, J. (1995). Neural mechanisms of selective visual attention. *Annual Reviews in Neuroscience, 18,* 193–222.

Duncan, J., Ward, R., & Shapiro, K. (1994). Direct measurement of attentional dwell time in human vision. *Nature, 369,* 313–315.

Elmore, L. C., Ji Ma, W., Magnotti, J. F., Leising, K. J., Passaro, A. D., Katz, J. S., et al. (2011). Visual short-term memory compared in rhesus monkeys and humans. *Current Biology, 21,* 975–979.

Enns, J. T., & Di Lollo, V. (1997). Object substitution: A new form of masking in unattended visual locations. *Psychological Science, 8,* 135.

Fisher, N. I. (1995). *Statistical analysis of circular data*. California: University Press.

Fougnie, D., Asplund, C. L., & Marois, R. (2010). What are the units of storage in visual working memory? *Journal of Vision, 10*(12):27, 1–11, http://www.journalofvision.org/content/10/12/27, doi:10.1167/10.12.27. [PubMed] [Article]

Gegenfurtner, K. R., & Sperling, G. (1993). Information transfer in iconic memory experiments. *Journal of Experimental Psychology: Human Perception and Performance, 19,* 845.

Gorgoraptis, N., Catalao, R. F. G., Bays, P. M., & Husain, M. (2011). Dynamic updating of working memory resources for visual objects. *Journal of Neuroscience, 31,* 8502.

Griffin, I. C., & Nobre, A. C. (2003). Orienting attention to locations in internal representations. *Journal of Cognitive Neuroscience, 15,* 1176–1194.

Hoffman, J. E. (1979). A two-stage model of visual search. *Attention, Perception, & Psychophysics, 25,* 319–327.

Jolicoeur, P., & Dell'Acqua, R. (1998). The demonstration of short-term consolidation. *Cognitive Psychology, 36,* 138.

Kastner, S., & Ungerleider, L. G. (2000). Mechanisms of visual attention in the human cortex. *Annual Reviews in Neuroscience, 23,* 315–341.

Lakha, L., & Wright, M. J. (2004). Capacity limitations of visual memory in two-interval comparison of Gabor arrays. *Vision Research, 44,* 1707–1716.

Lamy, D., & Zoaris, L. (2009). Task-irrelevant stimulus salience affects visual search. *Vision Research, 49,* 1472–1480.

Logie, R. H. (1995). *Visuo-spatial working memory*. Hove: Lawrence Erlbaum Associates.

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

Miller, E. K., Erickson, C. A., & Desimone, R. (1996). Neural mechanisms of visual working memory in

prefrontal cortex of the macaque. *Journal of Neuroscience, 16,* 5154.

Palmer, J. (1990). Attentional limits on the perception and memory of visual information. *Journal of Experimental Psychology: Human Perception and Performance, 16,* 332–350.

Palva, S., Kulashekhar, S., Hämäläinen, M., & Palva, J. M. (2011). Localization of cortical phase and amplitude dynamics during visual working memory encoding and retention. *Journal of Neuroscience, 31,* 5013–5025.

Pashler, H. (1988). Familiarity and visual change detection. *Perception & Psychophysics, 44,* 369–378.

Phillips, W. A. (1974). On the distinction between sensory storage and short-term visual memory. *Perception & Psychophysics, 16,* 283–290.

Posner, M. I. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology, 32,* 3–25.

Robertson, L. C. (2003). Binding, spatial attention and perceptual awareness. *Nature Reviews Neuroscience, 4,* 93–102.

Shibuya, H., & Bundesen, C. (1988). Visual selection from multielement displays: Measuring and modeling effects of exposure duration. *Journal of Experimental Psychology: Human Perception and Performance, 14,* 591–600.

Sperling, G. (1960). The information available in brief presentations. *Psychological Monographs, 74,* 1–29.

Theeuwes, J., & Burger, R. (1998). Attentional control during visual search: The effect of irrelevant single-tons. *Journal of Experimental Psychology: Human Perception and Performance, 24,* 1342–1353.

Treisman, A. (1998). Feature binding, attention and object perception. *Philosophical Transactions of the Royal Society of London B: Biological Sciences, 353,* 1295–1306.

Vogel, E. K., Woodman, G. F., & Luck, S. J. (2006). The time course of consolidation in visual working memory. *Journal of Experimental Psychology: Human Perception and Performance, 32,* 1436–1451.

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*(12):11, 1120–1135, http://www.journalofvision.org/content/4/12/11, doi:10.1167/4.12.11. [PubMed] [Article]

Wolfe, J. M. (1994). Guided Search 2.0: A revised model of guided search. *Psychonomic Bulletin & Review, 1,* 202–238.

Wolfe, J. M., & Cave, K. R. (1999). The psychophysical evidence for a binding problem in human vision. *Neuron, 24,* 11–17.

Woodman, G. F., & Vogel, E. K. (2005). Fractionating working memory. *Psychological Science, 16,* 106–113.

Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature, 453,* 233–235.