Little Support for Discrete Item Limits in Visual Working Memory

Klaus Oberauer
University of Zurich

in press, Psychological Science

Klaus Oberauer, Department of Psychology, University of Zurich.

Email: k.oberauer@psychologie.uzh.ch

This research was supported by a grant from the Swiss National Science Foundation (project 100014_179002). All Data and modelling code are available on the OSF: osf.io/vwnbh.

Abstract

Some theorists argue that working memory is limited to a discrete number of items, and additional items are not encoded at all. Others assume that all items are represented with variable quality. Adam, Vogel, and Awh (2017) presented evidence supporting the item-limit hypothesis: Participants reproduced visual features of up to six items in a self-chosen order. After the third or fourth response, error distributions were indistinguishable from guessing. I present four experiments with young adults (each N=24) testing the assumption that the brief, simultaneous display of visual arrays has led to failures of encoding in the experiments of Adam et al. (2017). Experiment 1 presented items slowly and sequentially. Experiment 2 presented them simultaneously but longer than in the experiments of Adam et al. (2017). Experiments 3 and 4 exactly replicated one original experiment. There was no evidence for an encoding limit. However, all four experiments failed to replicate the evidence for guessing-like error distributions. Modelling data from individuals revealed a mixture of some who do and others who don't produce guessing-like distributions. This heterogeneity increases the credibility of an alternative to the item-limit hypothesis: Some individuals decide to guess on hard trials even when they have weak information in memory.

Key words: Working memory, continuous reproduction, capacity.

Statement of Relevance

The nature of the capacity limit of visual working memory has been a matter of intense debate over the past two decades. One important distinction is between “slot theories” that assume a discrete capacity limit on the number of items or objects that can be maintained in visual working memory on the one hand, and alternative theories in which it is assumed that all items of a memory set are encoded, though with variable degree of strength or precision. The study of Adam, Vogel, and Awh (2017) is an important piece of evidence in favor of slot theories. The present experiments follow up on their finding, and show little evidence for their key observation of essentially flat (uniform) error distributions, as they would be expected from guessing. As such, the present results shift the balance of evidence substantially in favor of the alternative theories.

Little Support for Discrete Item Limits in Visual Working Memory

Adam et al. (2017) published an influential article reporting results that they interpreted as “clear evidence for item limits in visual working memory”. They refer to the hypothesis that the capacity of visual working memory (VWM) is a discrete number of place holders (a.k.a. slots), each of which can hold one item. The evidence comes from a whole-report continuous-reproduction test of VWM: Participants remembered arrays of up to six colors or orientations, and reported every item in that array by reproducing it on a continuous response scale, such as a color wheel. In the main experiments, participants were free to report the items in any order. Adam and colleagues found a strong preference of people to recall the items in descending order of how well they remembered them, so that the report error – the deviation of the reproduced feature from the true target feature – declined sharply from the first to the last reported item. For items reported after the first three, participants’ responses were statistically indistinguishable from pure guessing (i.e., a uniform distribution of responses over the entire circular response scale). Adam et al. (2017) concluded that participants were able to hold only three items in VWM, and had no memory information at all about any additional items. This conclusion is in line with the discrete-capacity hypothesis, but contrary to alternative theories, which share the assumption that all array items are encoded into VWM to some extent (Oberauer & Lin, 2017; Schneegans & Bays, 2017; van den Berg, Shin, Chou, George, & Ma, 2012).

The series of experiments I report here was motivated by the aim to identify boundary conditions for that finding. I started with the assumption that, by presenting arrays simultaneously and very briefly, Adam et al. (2017) might have made it hard for participants to even encode all items, especially at higher memory set sizes. Such encoding limits could have led to states of no information in VWM about some items, even without a discrete item limit on VWM. To test that possibility, Experiment 1 presented the array items slowly and sequentially, followed by whole report of all items in the participants’ chosen order. That experiment showed no evidence for zero-

information states in VWM. Experiment 2 took a step back towards the original experiments of Adam et al. (2017), presenting the arrays simultaneously, but long enough for encoding every item. Again, I found no evidence for zero-information states. Experiments 3 and 4 therefore were exact replications of the materials and procedure of one of the experiments of Adam et al. (2017); Experiment 3 tested the full range of set sizes investigated by Adam and colleagues, whereas Experiment 4 zoomed in on the largest and most diagnostic set size with a large number of trials to diagnose zero-information states in individual subjects. In these two experiments I was still not able to replicate the evidence that responses at later output positions were indistinguishable from guessing on the level of the entire sample, though a minority of subjects showed behavior compatible with pure guessing.

Methods

Participants. Students of the University of Zurich took part in each of the four experiments ($N = 23, 24, 23$, and 24 for Experiments 1, 2, 3, and 4, respectively) in exchange for partial course credit, or 45 Swiss Francs for 3 one-hour sessions (Experiments 1-3), or 70 Swiss Francs for 4 one-hour sessions (Experiment 4). The experiments were carried out in accordance with the guidelines of the Ethics Committee of the Faculty of Arts and Sciences at the University of Zurich.

Materials and Procedure. Materials and procedure were closely modelled after Experiment 1a of Adam et al. (2017), in which participants reproduced colors, and were free to choose the order of reporting all array items. The experiment was run with a program written in Matlab together with the Psychophysics toolbox 3 (Brainard, 1997), derived from the script of Adam et al. (2017).¹

In Experiments 1-3, set size of the arrays was 1, 2, 3, 4, or 6, with 90 trials per set size (30 in each session), presented in a random order. In Experiment 4 I used only set size 6 – the most diagnostic for zero-information states – to maximize the number of trials per condition to 500 (125

¹ I am grateful to Kirsten Adam, Ed Vogel, and Ed Awh for sharing their Matlab script, and their raw data, on the Open Science Framework.

in each session). The stimuli were colored squares of 50 x 50 pixels, presented on a grey background. For each array, their colors were drawn at random with replacement from 360 colors in a color circle in the CIE Lab color space centered at $L = 54$, $a = 18$, and $b = -8$. The stimuli were placed on equidistant locations on an invisible circle centered on the screen center, with a radius of 110 pixels.

In Experiment 1, the stimuli were presented sequentially for 500 ms each, starting at the top of the invisible circle and proceeding in clockwise order. In Experiment 2, the stimuli were presented simultaneously for 250 ms times the set size. In Experiments 3 and 4, stimuli were presented simultaneously for 150 ms regardless of set size, as in Adam et al. (2017).

After a 1 s retention interval during which the screen went blank, the color wheel was displayed in a new random rotation on each trial, together with dark-grey squares in the locations of the array items. Participants were instructed to report each item by choosing its color from the color wheel with a mouse click, and to do so in any order they liked. For each response they had to first select the item to be reported by clicking on its place holder, then click on the color for that item. In this way they had to report the colors of all items in the array. Participants were instructed to click on the colors they believed they remembered – even if only vaguely - with the left mouse button, but when they thought they had no clue and were merely guessing, they were to click on it with the right mouse button.

To prevent participants from encoding the stimuli verbally, in Experiments 1 and 2 we asked them to repeat aloud “ba bi bu” during presentation at a rate of one syllable every 500 ms (Experiment 1) or one syllable every 250 ms (Experiment 2). The speech rate was indicated by presenting the three syllables at the required rate on the screen before each trial. Participants’ speech was recorded, and spot-checked for compliance.

No participants or data were excluded from analysis. All raw data and modelling code (R and Jags) are available on the Open Science Framework (OSF): osf.io/vwnbh.

Results

I first report results pertaining to the question whether the data revealed evidence for zero-information states in VWM, followed by an exploratory analysis of people's meta-cognitive accuracy. An analysis of the serial-position data from Experiment 1 can be found in the Supplementary Material.

Evidence for Memory States without Information?

I computed the error of each response as the absolute angular deviation of the selected color from the true color in the color wheel. Figure 1 shows the distribution of errors for each set size and output position from Experiment 1. Figures 2 to 4 show the analogous set of error distributions for Experiments 2 to 4, respectively. Whereas Adam et al. (2017) have observed virtually flat (uniform) distributions of errors at output positions larger than 3, this was not the case here: At every set size and output position the error distributions peaked around zero, showing that participants had some information about the item's color at least on a subset of trials.

Figure 1. Error distributions from Experiment 1. *SS* = set size, *Out* = output position.

Figure 2. Error distributions from Experiment 2. *SS* = set size, *Out* = output position.

Figure 3. Error distributions from Experiment 3. SS = set size, Out = output position.

Figure 4. Error distributions from Experiment 4. Out = output position.

To estimate the contribution of memory information to the error responses, I fit a mixture model to the data of each experiment, separately for each set size and output position. The mixture model is a measurement model for continuous reproduction introduced by Zhang and Luck (2008) to measure the number of items for which people could, on average, give a response informed by memory of the target. The three-parameter version that I used was subsequently proposed by Bays, Catalao, and Husain (2009). It describes the distribution of responses as a mixture of three distributions: (1) A von-Mises distribution – a normal distribution on the circle – centered on the target item’s true color, (2) a set of von-Mises distribution, each of which is centered on the color of one of the non-target items in the array, and (3) a uniform distribution. The target-centered distribution (1) reflects memory of the to-be-reported item’s feature with a precision given by the dispersion parameter of the von-Mises distribution. The non-target centered distributions (2) reflect binding errors in VWM: People confuse the color of the item in the to-be-reported location with the color of another array item in another location, reflecting a failure of bindings between colors and their locations in memory. The uniform distribution (3) can be interpreted as reflecting guessing when people have no information about the to-be-reported item. The model has three free parameters: $P(mem)$, The proportion of responses attributed to the target-centered distribution (1); $P(non-target)$, the proportion of responses attributed to non-targets (2), and κ , the precision of the

von-Mises distributions. The first two parameters jointly determine the proportion of responses attributed to guessing (3), $P(\text{uniform}) = 1 - P(\text{mem}) - P(\text{non-target})$.

I used the Bayesian hierarchical implementation of the mixture model (Oberauer, Stoneking, Wabersich, & Lin, 2017). The hierarchical model version describes each person's parameter values as coming from a population distribution. Fitting the model yields parameter estimates for the population (i.e., mean and standard deviation of the population distributions) as well as for each person. Because the person parameter estimates are constrained by the population-level parameters, their estimates are informed not only by the data of that person but also – indirectly – by the data of all other participants, yielding more stable person-parameter estimates than from models fit to each individual separately. The Bayesian implementation yields posterior probability distributions for all parameters, which provide information not only about their location (e.g., the posterior mean or median) but also about the degree of (un-)certainty of the estimation (i.e., the dispersion of the posteriors). I applied the model with the same priors as in Oberauer et al. (2017).

Figure 5 shows the parameter estimates from the mixture model across output positions of set size 6, for five experiments: The four experiments presented here, and the original Experiment 1a in Adam et al. (2017). Each data point is the mode of the posterior distribution of the population-level mean parameter (i.e., the parameter value with the highest posterior probability), and the error bars represent the 95% highest-density interval (i.e., the smallest interval in which the true parameter value lies with 95% posterior probability).

For the original experiment of Adam et al. (2017), the posterior of $P(\text{mem})$ dropped to values close to zero at the later output positions. As the proportion of non-target responses, after initially rising up to output position 3, also dropped to zero, this implies that the proportion of responses attributed to the uniform mixture component approached one, as shown in the lower-right panel of Figure 5. The highest-density interval of $P(\text{uniform})$ also reached one, meaning that the assumption that all responses at the last one or two output positions come from a zero-information state is fairly

credible. By contrast, for the present experiments the posteriors of $P(\text{uniform})$ remained clearly below one; their highest-density intervals all remained below 0.9. This means that the proportion of responses that can be attributed to zero-information states most likely does not exceed 0.9 at any output position. In other words: At every output position of the set-size 6 condition of the present experiments, there were at least some trials that reflect information about the array rather than guessing.

Figure 5: Modes of the posterior distributions of the population-level parameters of the 3-parameter mixture model. Error bars are 95% highest-density intervals.

Are There Individuals with Zero-Information States?

Whereas I found no evidence for zero-memory states in the population-level parameters, these parameters only reflect the average of the population. There remains the possibility that some individuals had zero-information states. I tried to identify such individuals through two converging approaches. First, I fit the 2-parameter mixture model (Zhang & Luck, 2008) to the data of each subject separately, for all output positions of set sizes 4 and 6. I fit the model twice to each data set, once with the Nelder-Mead algorithm as implemented in the *dfoptim* package, and once with the L-BFGS-B algorithm in the *optim* function of R, and kept the better of the two fits, that is, the one with the smaller deviance = $-2 \cdot \log(\text{likelihood})$. For each subject and condition, I compared the fit of the mixture model to that of a uniform model that reflects 100% random guesses. For model comparison I used the BIC, a fit indicator combining the deviance with a penalty for the extra flexibility that the mixture model has due to its 2 free parameters, whereas the uniform model has none. When the BIC favored the uniform model, I regarded the participant to have had a zero-information state for the given set size and output position. Table 1 shows the number of participants (out of 24) in each experiment that was identified in this way as having a consistent zero-information state for a given design cell.

Table 1: Number of participants for which a uniform model fits the error distributions better than a 2-parameter mixture model, for each set-size and output-position combination.

Exp.	[4,1]	[4,2]	[4,3]	[4,4]	[6,1]	[6,2]	[6,3]	[6,4]	[6,5]	[6,6]
1	0	0	0	0	0	0	0	2	9	10
2	0	0	0	0	0	0	0	1	5	10
3	0	0	0	0	0	0	0	3	9	15
4	--	--	--	--	0	0	1	3	6	11
Adam et al. (2017)	0	0	3	7	0	0	2	14	22	20

The data in Table 1 suggest that about half of the participants in each experiment had zero-information states for at least one combination of set size and output position. However, we need to keep in mind that these results rest on model fits to the data of individual subjects, which are more prone to distortion by measurement noise than the estimates from a hierarchical model.

To mitigate that limitation, in the second approach I used again a Bayesian hierarchical modelling framework. I extended the 2-parameter mixture model (Zhang & Luck, 2008) by a second mixture parameter for the probability that a subject is in a consistent zero-information state for an entire condition of an experiment. This double-mixture model therefore has mixture parameters on two levels, the subject level and the trial level: On a subject level, the distribution of subjects is modelled as a mixture of two populations: A population of subjects who always respond in a guessing-like manner, producing uniformly distributed errors, and a population of subjects who behave according to the Zhang & Luck mixture model. On the trial level, subjects from the mixture-model population produce an error distribution that is a mixture of trials reflecting memory information, and other trials reflecting guessing.

The hierarchical double-mixture model (HDMM) predicts each response $\hat{x}_{i,j}$ of participant j in trial i as distributed according to a von Mises distribution:

$$\hat{x}_{i,j} \sim VM(m_{i,j}; x_i z_{i,j} \kappa_j)$$

The mean of the von Mises, $m_{i,j}$, is the true feature of the target item. Conveniently, the von Mises distribution with a precision of 0 is a uniform distribution on the circle. Therefore, we can multiply the precision parameter κ_j with two binary variables, x_j and $z_{i,j}$, both drawn from a Bernoulli distribution:

$$\begin{aligned} x_i &\sim \text{Bern}(P_{\text{Mix}}) \\ z_{i,j} &\sim \text{Bern}(Pm_j) \end{aligned}$$

The first, x_j expresses whether subject j is from the mixture-model population ($x_j = 1$) or from the consistently-guessing population ($x_j = 0$). The probability that a subject comes from the mixture-model distribution is parameter P_{Mix} . The second binary variable, $z_{i,j}$, expresses whether a mixture-model subject's response on trial i actually comes from memory ($z_{i,j} = 1$) or from guessing ($z_{i,j} = 0$). The probability that person j responds based on memory is given by parameter Pm_j . If and only if both binary variables are 1, the response is modelled as coming from a von-Mises distribution with precision κ_j ; otherwise it is drawn from a uniform distribution. The prior for P_{Mix} was a uniform distribution in the range [0, 1].

The subject-level parameters are modeled as draws from a population distribution:

$$\begin{aligned} Pm_j &\sim \text{Beta}(A_{Pm}, B_{Pm}) \\ \kappa_j &\sim \text{Gamma}(S_\kappa, R_\kappa) \end{aligned}$$

The priors for these two distributions are the same as in the standard mixture model (Oberauer et al., 2017), with one exception: I used $\text{Gamma}(5, 0.2)$ as the prior for the mean of precision, which assigns low prior probability to small precision values. This represents the assumption that persons from the mixture-model population have a memory precision that is clearly distinguishable from guessing.

I applied the HDMM to the data of the last output position of set size 6 in each experiment, as that is the condition in which the first approach detected the most subjects with uniform error

distributions. Figure 6 shows the posterior distribution of P_{Mix} for the 5 experiments. For the present experiments, high values of P_{Mix} are most credible, implying that the probability of encountering a person with consistent guessing behavior in the population is low. For the data of Adam et al. (2017), by contrast, the posterior of P_{Mix} was poorly located. This means that in light of the data, nearly every value of P_{Mix} is credible. Hence, these data of Adam et al. (2017) are compatible with the conjecture that most or all subjects are in consistent zero-information states, but they are equally compatible with the assumption that none of them are.

Figure 6: Posterior probability distribution of the P_{Mix} parameter from the double-mixture model, which reflects the probability that a subject comes from the hypothetical population of subjects who have some memory information at output position 6 of the set-size 6 condition.

For each subject, the proportion of posterior samples of x_i that are 0 provides an estimate of posterior probability that the subject belongs to the consistent-guessing population. Figure 7 plots the two indicators of individual zero-information states against each other, the BIC difference from the maximum-likelihood model comparison, and the posterior probability of a consistent zero-information state from the Bayesian HDMM. The two indicators show a nearly perfect rank-order correlation in each experiment, but the posterior probability of $x_i = 0$ from the HDMM rarely exceeded 0.5 even for subjects for which the BIC difference was in favor of the uniform model over the mixture model.

To further check the validity of the classification of error distributions as coming from zero-information states, I selected the data of set size 6, output position 6, from those participants whose data in that design cell were fit better by a uniform model than the mixture model according to BIC, and plotted their error distributions in Figure 8. Whereas the error distribution of the data from Adam et al. (2017) plausibly come from a uniform distribution, those from the present experiments show a peak around zero, indicating that at least some of the participants identified as having had zero-information states had some memory for the true feature on at least some of the trials.

Figure 7: Correspondence of Indicators of Individual Zero-Information States. Positive BIC differences from the maximum-likelihood model comparison reflect evidence for the uniform model over the 2-parameter mixture model. Each data point comes from one subject in the set-size 6, output-position 6 condition. One data point with BIC difference of -203 was omitted from Experiment 4.

Figure 8: Error distributions from set-size 6, output position 6, for the subset of participants who were identified as having had a zero-information state by the individual maximum-likelihood fits of the 2-parameter mixture model.

Meta-Cognitive Accuracy

Participants were asked to respond with the left mouse button when they remembered the color, and with the right mouse button when they thought they were only guessing. Figure 9 plots error distributions from set size 6 – the condition with the most even split between self-reported memory-based and guessing responses. In all four experiments, the error distributions from self-reported guesses were approximately uniform.

Apparently, participants had good meta-cognitive insight into when their responses were informed by memory and when they were not – at least they were very good at identifying the latter situation. At the same time, the error distributions from self-reported memory-based responses still showed a substantial number of very large errors, reflected in the broad tails of the distributions. This could be because, as participants have been instructed to use the right mouse button conservatively, they used the left mouse button in case of doubt. In this interpretation, the decision to use the left or the right mouse button is made as a signal-detection process on a continuously varying degree of confidence, with a conservative criterion.

Figure 9: Error distributions from set size 6 for responses with the left mouse button (Guessing=0) and with the right mouse button (Guessing = 1) in Experiments 1 (top), 2 (middle) and 3 (bottom).

Alternatively, self-reported guesses could reflect a qualitatively different meta-cognitive state from self-reported memory-based responses. Self-reported guesses could reflect the trials on which participants made a deliberate decision to guess, perhaps because they found the attempt to remember the color hopeless, or not worth the effort of a lengthy retrieval attempt. By contrast, responses with the left mouse button reflect genuine memory-based responses, some of which happen to reflect extremely poor, noisy memory, leading to very large errors. Future research measuring people's confidence with a more fine-grained, approximately continuous scale could perhaps help to adjudicate between these possibilities.

People's highly accurate meta-cognitive judgments agree with earlier findings showing a strong relation between confidence and precision of reproduction (Honig, Ma, & Fournie, 2020; Mitchell & Cusack, 2018; Rademaker, Tredway, & Tong, 2012). Moreover, they converge with the observation that the present participants – as those of Adam et al. (2017) – were very good at ordering their responses by reproduction accuracy.

Discussion

The initial motivation of the present series of experiments was to investigate the possibility that the zero-information states that Adam et al. (2017) have identified reflected limitations of encoding rather than limitations of working memory. The results provided evidence against that assumption: In Experiment 1, colors were presented sequentially at a relatively slow pace, giving participants ample opportunity to attend each of them. By contrast, Experiments 3 and 4 replicated the brief, simultaneous presentation of arrays as in the original study of Adam et al. (2017), which could have caused limitations of encoding. Neither overall performance, nor the evidence for zero-information states in individual conditions or participants, differed between these experiments. Therefore, presenting arrays of colors simultaneously for merely 150 ms does not impose a noticeable limit on encoding. This conclusion converges with those from previous studies that found only small, and inconsistent, differences between simultaneous and sequential presentation of simple visual stimuli (Gorgoraptis, Catalao, Bays, & Husain, 2011; Liu & Becker, 2013; Miller, Becker,

& Liu, 2014), and little gain in memory performance from exposure times longer than a few 100 ms (Bays, Gorgoraptis, Wee, Marshall, & Husain, 2011; Vogel, Woodman, & Luck, 2006).

Nevertheless, none of the present experiments replicated the observation of uniformly distributed error distributions at the group level, which reflect essentially no information from memory. Therefore, the “clear evidence for item limits” that Adam et al. (2017) have identified in their data was not found here. This result converges with a recent study by Bruning and Lewis-Peacock (2020), who tested continuous reproduction of six locations in a free-ordered whole-report procedure, and found error distributions centered on the true location at all six output positions. An analysis on the level of subjects revealed individual differences: About half the subjects in the present experiments showed evidence consistent with zero-information states; the other half did not; in the data of Adam and colleagues, only a small minority was identified as not showing zero-information states. Therefore, the present results do not deviate qualitatively from those of Adam and colleagues but quantitatively, reflecting a larger proportion of subjects without zero-information states. As the samples came from different student populations, it is plausible that they reflect a different mix of individual strategies.

Whereas the present findings do not lend support to the hypothesis of discrete item limits in VWM, they do not provide clear evidence against it either, for the following reason: The prediction of uniformly distributed error distributions at large set sizes and late output positions follows from the conjunction of two assumptions: One is that VWM is limited to a discrete number of items, and the other is that participants very consistently start reproducing the features of items they have in VWM, followed by guessing on those items that they don’t have. If one of these assumptions is wrong, the prediction no longer holds. It is possible that the second assumption is wrong for the present experiments: Participants might have been less consistent in recalling the items they know first. For instance, a person with an item limit of 3 or 4 could – at least on a substantial proportion of trials – give their report of a six-item array in a left-to-right, top-to-bottom order, regardless of which

items they hold in working memory. In this way, some responses even at the last output position would reflect memory, whereas some earlier responses would reflect guesses. To conclude, the design of the present experiments, and those of Adam et al. (2017), is suited only for providing evidence for the existence of zero-information states in VWM, not against it.

That said, the existence of uniform error distributions in Adam et al. (2017) is not unambiguous evidence for discrete item limits either. The assumption of a discrete capacity limit entails that, if set size exceeds capacity, people must have a zero-information state about some of the items, as reflected in a uniform error distribution on the circle. The converse, however, does not hold, because there are alternative explanations for uniform error distributions. One is that they reflect the tail of a broad distribution of memory precisions (van den Berg et al., 2012): Precision fluctuates widely between and within trials, and some items are encoded with such low precision that their reproduction is not measurably better than guessing. A recent model comparison has shown a better quantitative account of the Adam et al. data by a variable-precision model than by a discrete-capacity model (Schneegans, Taylor, & Bays, 2020). This explanation would go along with the assumption that confidence varies continuously as a function of precision (van den Berg, Yoo, & Ma, 2017), and people make their decisions whether or not to report that they were guessing by placing a criterion on that continuum. Adam et al. (2017) argued that this is a theoretically unattractive proposition: A variable-precision model has to allow precision to vary so much that, in some conditions, the precision of items is indistinguishable from zero. Such a model would exploit the flexibility of a precision distribution that can be stretched out freely to mimic the behavior of a discrete-capacity model.

Another alternative explanation is that people have representations of non-negligible precision of all array items, but choose to guess on some responses nonetheless (Oberauer & Lin, 2017). Although doing so would mean to sacrifice available memory information, it could still be rational because retrieving that information might not be worth the time investment when the

information is difficult to retrieve, and likely to be poor. In this account, there is a qualitative distinction between memory-based responses and guesses, but that distinction pertains not to the state of memory, but to the person's meta-cognitive decision on whether to use memory. Applied to the findings of Adam et al. (2017), this explanation is not very plausible, as one had to assume that nearly every person decided to guess on nearly every trial at output positions 4 and higher. Such a high agreement on a strategic decision that involves sacrificing some degree of response accuracy is unlikely. The present finding that at best about half of the subjects produced response distributions indistinguishable from guessing substantially mitigates that problem: We only have to assume that some people decide to forfeit a small gain in accuracy by not retrieving weak memories. In light of individual differences in achievement motivation and in strategy choices, this is a fairly credible assumption.

To conclude, neither is the existence of uniform error distributions decisive evidence for a discrete capacity limit on items in WM, nor is their absence decisive evidence against it. Nevertheless, the demonstration of uniform error distributions in the experiments of Adam et al. (2017) precisely in those conditions in which one would expect them on the basis of a capacity limit to about 3-4 items increased the credibility of that hypothesis. Expressed in Bayesian terms: The likelihood of such a finding is high under the assumption of a discrete capacity limit, and lower – though not zero – for alternative models. It is lower because alternative models had to invoke auxiliary assumptions with low prior probability. Thereby, the discrete-capacity hypothesis gained in posterior probability relative to competing accounts of WM capacity limitations. By the same rationale, the present findings shift the balance of evidence more in favor of the hypothesis that all items of a visual array are represented in VWM with some non-negligible strength and precision, because the necessary assumptions in alternative models now have higher prior probability.

References

- Adam, K. C. S., Vogel, E. K., & Awh, E. (2017). Clear evidence for item limits in visual working memory. *Cognitive Psychology*, 97, 79-97. doi:10.1016/j.cogpsych.2017.07.001
- 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, 1-11.
- Bays, P. M., Gorgoraptis, N., Wee, N., Marshall, L., & Husain, M. (2011). Temporal dynamics of encoding, storage, and reallocation of visual working memory. *Journal of Vision*, 11, 1-15.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, 10, 433-436.
- Bruning, A. L., & Lewis-Peacock, J. A. (2020). Long-term memory guides resource allocation in working memory. *Scientific Reports*, 10(1), 22161. doi:10.1038/s41598-020-79108-1
- Cortis, C., Dent, K., Kennett, S., & Ward, G. (2015). First things first: Similar list length and output order effects for verbal and nonverbal stimuli. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 41, 1179-1214. doi:10.1037/xlm0000086
- 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, 8502-8511.
- Grenfell-Essam, R., Ward, G., & Tan, L. (2013). The role of rehearsal on the output order of immediate free recall of short and long lists. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 39, 317-347. doi:10.1037/a0028974
- Honig, M., Ma, W. J., & Fougner, D. (2020). Humans incorporate trial-to-trial working memory uncertainty into rewarded decisions. *Proceedings of the National Academy of Sciences*, 117(15), 8391. doi:10.1073/pnas.1918143117
- Liu, T., & Becker, M. W. (2013). Serial consolidation of orientation information into visual short-term memory. *Psychological Science*. doi:10.1177/0956797612464381
- Miller, J. R., Becker, M. W., & Liu, T. (2014). The bandwidth of consolidation into visual short-term memory depends on the visual feature. *Visual Cognition*, 22, 920-947. doi:10.1080/13506285.2014.936923
- Mitchell, D. J., & Cusack, R. (2018). Visual short-term memory through the lifespan: Preserved benefits of context and metacognition. *Psychology and Aging*, 33(5), 841-854. doi:10.1037/pag0000265
- Oberauer, K. (2003). Understanding serial position curves in short-term recognition and recall. *Journal of Memory and Language*, 49(4), 469-483.
- Oberauer, K., Lewandowsky, S., Awh, E., Brown, G. D. A., Conway, A. R. A., Cowan, N., . . . Ward, G. (2018). Benchmarks for models of short-term and working memory. *Psychological Bulletin*, 144, 885-958. doi:10.1037/bul0000153
- Oberauer, K., & Lin, H.-Y. (2017). An interference model of visual working memory. *Psychological Review*, 124, 21-59.
- Oberauer, K., Stoneking, C., Wabersich, D., & Lin, H.-Y. (2017). Hierarchical Bayesian measurement models for continuous reproduction of visual features from working memory. *Journal of Vision*, 17, 1-27. doi:10.1167/17.5.11
- Peteranderl, S., & Oberauer, K. (2018). Serial recall of colors: Two models of memory for serial order applied to continuous visual stimuli. *Memory & Cognition*, 46, 1-16. doi:10.3758/s13421-017-0741-0
- Rademaker, R. L., Tredway, C. H., & Tong, F. (2012). Introspective judgments predict the precision and likelihood of successful maintenance of visual working memory. *Journal of Vision*, 12, 1-13. doi:10.1167/12.13.21
- Schneegans, S., & Bays, P. M. (2017). Neural architecture for feature binding in visual working memory. *Journal of Neuroscience*, 37, 3913-3925. doi:10.1523/JNEUROSCI.3493-16.2017
- Schneegans, S., Taylor, R., & Bays, P. M. (2020). Stochastic sampling provides a unifying account of visual working memory limits. *Proceedings of the National Academy of Sciences*, 117(34), 20959. doi:10.1073/pnas.2004306117

- van den Berg, R., Shin, H., Chou, W.-C., George, R., & Ma, W. J. (2012). Variability in encoding precision accounts for visual short-term memory limitations. *Proceedings of the National Academy of Sciences*, *109*, 8780-8785.
- van den Berg, R., Yoo, A. H., & Ma, W. J. (2017). Fechner's law in metacognition: a quantitative model of visual working memory confidence. *Psychological Review*, *124*, 197-214. doi:10.1037/rev0000060
- 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.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233-236.

Supplementary Material

Serial-Position Effects

Experiment 1, in which items were presented sequentially, offers the opportunity to analyze reproduction errors and self-reported guesses as a function of the items' serial position in the presentation sequence (i.e., input position), as well as their position in the recall sequence (i.e., output position). We can also investigate how people's self-chosen order of recall depends on the order of presentation.

Participants showed a strong tendency to report the items in the order in which they were presented, and that tendency was strongest for the smaller set sizes. Table 2 shows the frequencies of the 7 most popular output orders at each set size

Table 2: Frequencies of the 7 most popular output orders in Experiment 1

Set Size 2 Order	Frequency	Set Size 3 Order	Frequency	Set Size 4 Order	Frequency	Set Size 6 Order	Frequency
12	1829	123	1421	1234	919	123456	486
21	119	132	221	1423	145	612345	124
		321	122	1324	130	654321	11
		312	75	1243	107	162345	68
		231	63	4321	99	561234	61
		213	46	4123	92	651234	51
				1432	69	165432	40

Note: Orders are given by listing the input positions of reported items in the order in which they were reported.

This observation matches findings from free-recall tasks with words, and with spatial locations (Cortis, Dent, Kennett, & Ward, 2015; Grenfell-Essam, Ward, & Tan, 2013): For list lengths up to about 6 items, people have a preference for starting recall with the first-presented item, and tend to proceed in forward order. For that reason, the following analyses are split by trials in which participants' output order strictly matched the presentation order, and trials in which it did not.

Figure S1 shows the mean absolute error (on the left) and the proportion of trials on which participants indicated that they were merely guessing by responding with the right mouse button (right) as a function of input position (= output position) for those trials that were reported in strict forward order. Errors show the serial-position curves typical for forward serial recall of verbal and spatial lists (for a review see Oberauer et al., 2018), with an extended primacy effect and a small recency effect, observed here only for the largest set size. A similar pattern has been found with serial continuous-reproduction of colors (Peteranderl & Oberauer, 2018). The proportion of self-reported guesses increased steadily across output positions, but they remained a minority of trials even at the last output position.

Figure S1: Mean errors (absolute deviation of response from target color in degrees) and proportion of self-reported guessing responses in Experiment 1 for the subset of trials in which the colors were reported strictly in their order of presentation. SS = set size; the percentage in the legend is the percentage of trials reported in strict forward order.

The corresponding serial-position curves for the trials not recalled in forward order are plotted in Figure S2 as a function of input order in the top panels, and as a function of output order in the bottom panels. The serial-position curve of errors mirror those observed with verbal materials when items are probed in a random order to disentangle input and output position effects (Oberauer, 2003): There were approximately symmetric primacy and recency effects over input position, and a steady decline of performance across output positions. The same trends can be seen in the proportions of self-reported guesses.

Figure S2: Mean errors (absolute deviation of response from target color in degrees) and proportion of self-reported guessing responses in Experiment 1 for the subset of trials in which the colors were not reported in their order of presentation. SS = set size; the percentage in the legend is the percentage of trials reported in strict forward order.

Compared to the trials reported in forward order, here the proportion of self-reported guesses was substantially higher, whereas performance was at a comparable level. This could be a hint that people chose to recall in forward order to the extent that they were confident to be able to

recall all items – in agreement with the observation that people’s inclination to recall a list in forward order decreases with larger set size, and also when recall is made harder by other experimental manipulations (Grenfell-Essam et al., 2013).

Taken together, the effects of serial position on reproduction accuracy mirror findings from list-memory studies using other (primarily verbal) materials, supporting the assumption that primacy and recency effects, and the tendency to recall short lists in forward order in free-recall tests, reflect general principles of memory.