

A Neural Measure of Precision in Visual Working Memory

Edward F. Ester¹, David E. Anderson², John T. Serences¹,
and Edward Awh²

Abstract

Recent studies suggest that the temporary storage of visual detail in working memory is mediated by sensory recruitment or sustained patterns of stimulus-specific activation within feature-selective regions of visual cortex. According to a strong version of this hypothesis, the relative “quality” of these patterns should determine the clarity of an individual’s memory. Here, we provide a direct test of this claim. We used fMRI and a forward encoding model to characterize population-level orientation-selective responses in visual cortex while human participants held an oriented grating in memory. This analysis, which enables a

precise quantitative description of multivoxel, population-level activity measured during working memory storage, revealed graded response profiles whose amplitudes were greatest for the remembered orientation and fell monotonically as the angular distance from this orientation increased. Moreover, interparticipant differences in the dispersion—but not the amplitude—of these response profiles were strongly correlated with performance on a concurrent memory recall task. These findings provide important new evidence linking the precision of sustained population-level responses in visual cortex and memory acuity. ■

INTRODUCTION

Working memory (WM) enables the temporary storage of information in a readily accessible state. This system is critical for virtually all forms of “online” cognitive processing, as evidenced by robust correlations with measures of fluid intelligence and scholastic aptitude (e.g., Cowan et al., 2005). Research suggests that WM storage is mediated by a distributed network of prefrontal, parietal, and infero-temporal cortical areas (e.g., Xu & Chun, 2006; Todd & Marois, 2004; Miller, Erickson, & Desimone, 1996). However, many of these regions lack the fine-grained selectivity of early visual areas. Thus, recent investigations have begun to examine how humans store fine visual details over short intervals. An emerging perspective—informed by unit recordings in nonhuman primates (Pasternak & Greenlee, 2005; Super, Spekreijse, & Lamme, 2001) and human neuroimaging studies (Riggall & Postle, 2012; Ester, Serences, & Awh, 2009; Harrison & Tong, 2009; Serences, Ester, Vogel, & Awh, 2009)—is that this ability is mediated by sensory recruitment or sustained activity in sensory cortical areas that exhibit selectivity for the remembered feature (D’Esposito, 2007; Postle, 2006). For example, recent human neuroimaging studies (e.g., Harrison & Tong, 2009; Serences et al., 2009) have demonstrated that, when participants are asked to remember an orientation over a short delay, sustained patterns of activation in early visual cortex (e.g., V1–hV4) discriminate the specific orientation value being stored.

Sensory recruitment is thought to determine the relative precision mnemonic representations. By this account, individual differences in the quality of stimulus-specific patterns of activation observed during WM storage should predict the quality of observers’ memory representations. However, extant studies examining sensory recruitment have relied on relatively coarse “decoding” analyses that preclude formal quantitative descriptions of these patterns (see Serences & Saproo, 2012, for a detailed discussion of this issue). Moreover, these analyses are typically used to discriminate between highly dissimilar feature values (e.g., 25° vs. 115° orientations, Harrison & Tong, 2009; 45° vs. 135°, Ester et al., 2009; Serences et al., 2009) that are well above memory discrimination thresholds. Thus, it is unclear (1) whether sustained, stimulus-specific patterns of activation observed in visual cortex during WM storage are functionally linked to mnemonic acuity and (2) which aspect(s) of these profiles are responsible for this link.

Here, we report the results of two experiments designed to examine these questions. In Experiment 1, we used fMRI and a forward encoding model to quantify population-level orientation-selective responses in early visual areas (e.g., V1–V3v) while participants remembered a specific orientation value over a short delay. Critically, interparticipant differences in the dispersion—but not the amplitude—of these profiles were strongly correlated with performance on a concurrent memory recall task. In Experiment 2, we show that these orientation-selective response profiles are critically dependent on a participant’s intent to remember a specific feature value. Together,

¹University of California, San Diego, ²University of Oregon

these findings provide critical evidence for sensory recruitment models of memory by demonstrating that the relative “quality” of sustained stimulus-specific patterns observed in visual cortex during WM storage are functionally linked with memory acuity.

METHODS

Participants

Twenty-one students from the University of Oregon (ages 19–33 years) participated in Experiment 1 and nine students from the University of Oregon participated in Experiment 2. All participants reported normal or corrected-to-normal visual acuity, and all gave both written and oral informed consent. Participants were tested in a single 2-hr scanning session and compensated at a rate of \$25/hr. Data from one participant in Experiment 1 and one participant in Experiment 2 were discarded due to large motion artifacts. The data reported here reflect the remaining 20 and 8 participants, respectively.

Stimuli and Apparatus

Stimuli were generated in Matlab and back-projected onto a 32-cm screen via Psychophysics toolbox software (Brainard, 1997; Pelli, 1997). Participants were instructed to maintain fixation on a small black dot (subtending 0.2° from a viewing distance of 68 cm) present for the duration of the scan.

Experiment 1—WM Task

A representative trial is depicted in Figure 1. Each trial began with the presentation of a full-contrast square-wave grating (radius 8°) for 1 sec. This “sample” stimulus was rendered in one of eight orientations (0 – 157.5° in 22.5° increments; jittered on each trial by a ± 1 – 10° on each trial to discourage verbal coding) and flickered at 3 Hz (i.e., 167 msec on, 167 msec off). To attenuate the potency of retinal afterimages, the spatial phase of the grating was randomly assigned one of four equidistant values from 0 to 2π on each cycle. The sample stimulus was followed by a 12-sec delay interval and the presentation of a randomly oriented probe. Participants were given 3 sec to adjust the orientation of this probe to match that of

the remembered sample (using keys on a custom-made MR-compatible button box). Each trial was followed by a 2-sec ($n = 1$) or variable (3, 4, 6, or 7 sec; $n = 19$) blank intertrial interval (ITI). Each scan contained 16 trials and lasted either 288 (for the participant with fixed 2-sec ITIs) or 338 sec (for the remaining participants with variable ITIs). Each participant completed 7–12 scans (median = 9) as time permitted.


Experiment 2—Store versus Drop Memory Task

Experiment 2 was identical to Experiment 1. However, the sample was rendered in one of five orientations (0 – 144° in 36° increments; randomly jittered by ± 1 – 10° on each trial to discourage verbal coding) and followed by change in the color of the fixation point (to green or red; color mappings were counterbalanced across participants) that instructed participants to remember the sample’s orientation (“store” trials) or simply wait for the next trial begin (“drop” trials). This cue was present for the entire delay interval.

Retinotopic Mapping, Functional Localizer, and Voxel Selection

Retinotopic mapping data were acquired using a rotating checkerboard wedge subtending 45° and flickering at 8 Hz. This procedure was used to identify visual areas V1–V3v in each participant. Each participant also completed one scan (15 trials) in a functional localizer task. A full-contrast, phase-reversing (10 Hz) checkerboard stimulus (radius 8°) was presented for a total of 10 sec; participants were required to detect brief (50 msec) reductions in stimulus contrast that occurred at unpredictable intervals. Each trial was followed by a 10-sec fixation interval. To identify visually responsive voxels in visual areas V1–V3v, we constructed a general linear model with a single boxcar regressor denoting stimulus presence (i.e., on vs. off). The regressor was convolved with a gamma function to account for the assumed shape of the hemodynamic response. Voxels that showed a stronger response during epochs of stimulation (relative to fixation; thresholded at $p < .05$ using the false discovery rate algorithm provided in BrainVoyagerQX 1.9, Brain Innovations, The Netherlands) were used to define functional ROIs in visual areas V1–V3v.

Figure 1. Behavioral task. On each trial, participants were required to remember the orientation of a sample grating over a 12-sec interval and then adjust a randomly oriented probe to match the remembered sample.



fMRI Data Acquisition and Preprocessing

fMRI data were collected using a 3-T Siemens (Malvern, PA) Allegra system located at the University of Oregon. Anatomical images were acquired using a spoiled-gradient-recalled T1-weighted sequence that yielded images with 1-mm³ resolution. Whole-brain EPIs were acquired in 33 transverse slices using the following parameters: 3 mm in-plane resolution, 2000 msec repetition time, 30 msec echo time, 90° flip angle, 64 × 64 image matrix, 192 mm field of view, 3.5 mm slice thickness (no gap). EPIs were slice-time corrected, motion corrected (both within and between scans), and high-pass filtered (3 cycles/run). Preprocessing was performed using BrainVoyager QX 1.9 and custom routines written in Matlab.

Forward Encoding Model of Orientation Selectivity

To characterize orientation-specific responses in visual cortex during WM maintenance, we generated a set of orientation-selective response functions (channel response functions [CRFs]) using a forward encoding model of orientation selectivity. Our approach was similar to one described by Brouwer and Heeger (2009, 2011), and we therefore adopt their terminology and conventions throughout the manuscript. Briefly, the encoding model used here assumes that each fMRI voxel in visual cortex samples from a large number of orientation-selective neurons and that the response of any given voxel is proportional to the summed responses of all neurons in that voxel. Thus, one can characterize the orientation selectivity of a given voxel as a weighted sum of n orientation channels, each with an idealized tuning curve. Specifically, we modeled the response of each voxel using a basis set of eight half-rectified sinusoids (one per sample orientation) raised to the fifth power. These functions were chosen to approximate single-unit tuning profiles in primary visual cortex, where the $1/\sqrt{2}$ half-bandwidth of orientation-selective cells has been estimated at about 20° (although there is a considerable amount of variability in this estimate, e.g., Ringach, Shapley, & Hawken, 2002). However, qualitatively similar results were observed when basis functions were defined as sinusoids raised to the sixth or eighth power.

In the first phase of our analysis, we extracted the raw time series from each voxel in a given ROI during a time period extending from 8 to 12 sec following the start of each trial (qualitatively similar findings were obtained across a wide range of temporal windows and on a time point-by-time point basis; see below). Each time series was normalized on a scan-by-scan basis using a z transform and sorted into one of eight bins based on stimulus orientation. Data were subsequently pooled and averaged across corresponding ROIs in each visual area (i.e., left and right V1) as no hemispheric asymmetries were observed. Data from all but one scan were designated as a “training” set, and data from the remaining scan were designated as a

“test” set (partitioning the data in this manner ensures that the training and test sets are always independent).

In the second phase of the analysis, data from the training set were used to estimate weights on the hypothetical orientation channels separately for each voxel. Using the terminology of Brouwer and Heeger (2009, 2011), let m be the number of voxels in an ROI, k be the number of hypothetical orientation channels, and n_1 and n_2 be the number of observations in the training and test sets, respectively. The channel weights (W , $m \times k$ can be derived via least-squares estimation),

$$W = D_1 C_1^T (C_1 C_1^T)^{-1} \quad (1)$$

where D_1 ($m \times n_1$) is the training set and C_1 ($k \times n_1$) is the basis set discussed above. In the second phase of the analysis, we estimated channel responses (C_2 , $k \times n_2$) using these weights and data from the test set (D_2 , $m \times n_2$),

$$C_2 = (W^T W)^{-1} W^T D_2 \quad (2)$$

Importantly, the responses in C_2 are scaled by the amplitudes of the sinusoids within the basis set (all of which were set to 1). Thus, units of response are arbitrary.

The columns in C_2 were then circularly shifted so that the channel tuned to the stimulus presented on each trial was positioned in the center of orientation space, thereby aligning the estimated channel responses to a common center (i.e., 0°). This analysis was repeated iteratively until all scans had served as the test set and the results were averaged. Finally, to characterize the amplitude and dispersion of these functions, each participant’s response profile was fit with a Gaussian function of the form,

$$f(x) = ae^{\sigma \cos(x-\mu)} + b \quad (3)$$

where μ and σ correspond to the mean and dispersion (i.e., the inverse of the standard deviation) of the distribution, respectively. a and b are scaling factors that correspond to the amplitude and baseline of the observed function, respectively. a was taken as an estimate of response amplitude, and σ was taken as an estimate of dispersion.

Forward Encoding Model—Experiment 2

A qualitatively similar encoding model was used to characterize orientation-selective responses during “store” and “drop” trials in Experiment 2. Data from all but one scan were designated as a “training” set and used to estimate weights on five hypothetical orientation channels (sinusoids raised to the fifth power). Critically, the training set contained data from both “store” and “drop” trials; this ensured that estimated weights were unbiased across

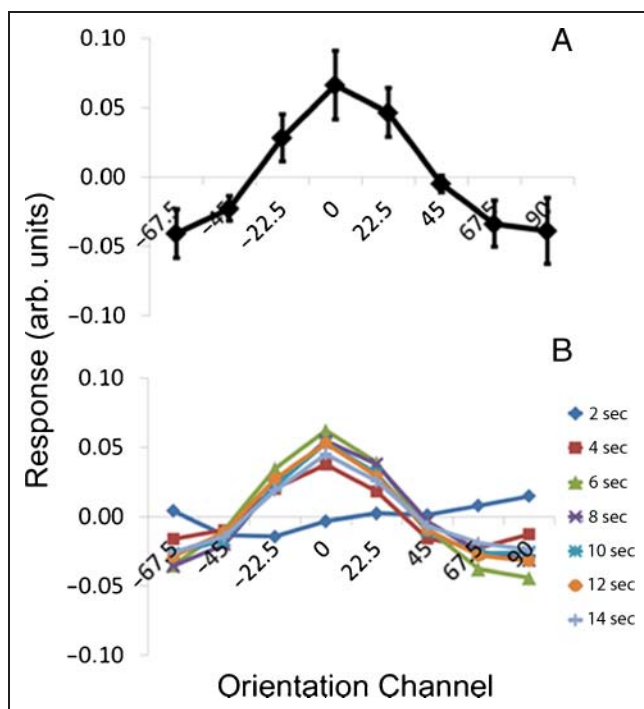


Figure 2. Characterizing stimulus-specific patterns observed in visual cortex during WM storage. (A) A forward encoding model (see Experimental Procedures) was used to generate an orientation-selective population response profile (CRF) based on patterns of activation observed in V1 and V2v during a period from 8 to 12 sec following the start of each trial. CRFs peaked in the channel preferring the remembered orientation and decreased monotonically as the angular distance from this orientation increased. Error bars are ± 1 SEM. (B) Each line depicts the mean CRF observed at a different point during the course of a trial. Error bars have been omitted for exposition.

conditions. Data from the remaining scan were designated as a “test” set and further partitioned into “store” and “drop” subsets. We then estimated channel responses for store and drop trials (separately) via Equation 2. Channel responses were circularly shifted to a common center, and the entire analysis was iterated until data from each scan had served as the test set. The results were averaged, yielding a single orientation-selective channel response profile for store and drop trials.

RESULTS

Experiment 1

For each participant, we computed a distribution of recall errors (i.e., the angular distance between the reported and sample orientation) across all trials. Response errors were tightly clustered around the sample orientation, with an average deviation of deviation of 7.4° from the sample orientation ($SEM = 0.4^\circ$). However, there was substantial variability in acuity (i.e., mean absolute recall error) across participants (range = 4.4° – 12.6°).

Next, we used a forward encoding model (Brouwer & Heeger, 2009, 2011) to generate population-level orientation-selective response functions based on pat-

terns of activation measured in visual cortex during the WM delay (see Methods). Briefly, we characterized the response of each visually responsive voxel in visual areas V1–V3v during the memory delay as a weighted sum of eight hypothetical orientation “channels” (one per sample orientation), each with an idealized tuning function. In the first phase of the analysis, data from all but one scan were used to estimate weights on each of the eight orientation channels separately for each voxel. In the second phase of the analysis, we estimated the response of each channel given the weights and data from the remaining scan. This procedure was iterated until all scans had served as the test set and the results were averaged (recall that data were partitioned into training and test sets on a scan-by-scan basis; thus, the two data sets were always statistically independent).

As shown in Figure 2A, this analysis revealed a graded orientation-selective profile of activation that peaked in the channel corresponding to the orientation stored in WM. The data shown in Figure 2A have been pooled across visual areas V1 and V2v; identical findings were observed when V1 and V2v were considered separately. We could not recover orderly tuning profiles for approximately 40% of subjects in area V3v. The functions depicted in Figure 2A correspond to data from a period of 8–12 sec following the start of the trial. However, tuning functions emerged approximately 4 sec following the start of each trial (i.e., 3 sec after sample offset) and persisted until the presentation of the probe stimulus (Figure 2B). Next, we quantified the amplitude and dispersion of each participant’s response profile by fitting his or her response profile with a Gaussian function; estimated values are plotted as a function of time in Figure 3. Once a robust channel response profile emerged (approximately 4 sec after the start of each trial; see Figure 2B), both dispersion and amplitude




Figure 3. Estimates of CRF amplitude and dispersion are largely unchanged during the memory delay. Mean amplitude (red; right ordinate) and dispersion (black; left ordinate) are plotted as a function of time (abscissa) relative to the start of each trial. Once CRFs began to emerge (approximately 4 sec following the start of each trial; see Figure 2B), estimates of dispersion and amplitude remained largely unchanged.

remained relatively constant over time: Separate one-way ANOVAs revealed no effect of time on estimates of dispersion (black line) or amplitude (red line; both p s > .40; Greenhouse–Geisser corrections were applied to account for violations of sphericity). Critically, however, individual differences in the dispersion of the observed response profiles were a robust predictor of participants' mean recall errors (Figure 4, top; $R^2 = 0.44$, $t(18) = 3.63$, $p < .01$) such that broader tuning profiles were associated with greater error.¹ In fact, this relationship was robust across nearly the entire delay interval (Figure 4, bottom) as well as when each visual area (i.e., V1 and V2v) was considered independently (using data from a period of 8–12 sec following the start of each trial; $R^2 = 0.27$ and 0.30 , respectively; $p < .05$). No relationship between dispersion and recall performance was observed in visual area V3v ($R^2 = 0.003$). Conversely, the amplitudes of the observed response profiles were uncorrelated with

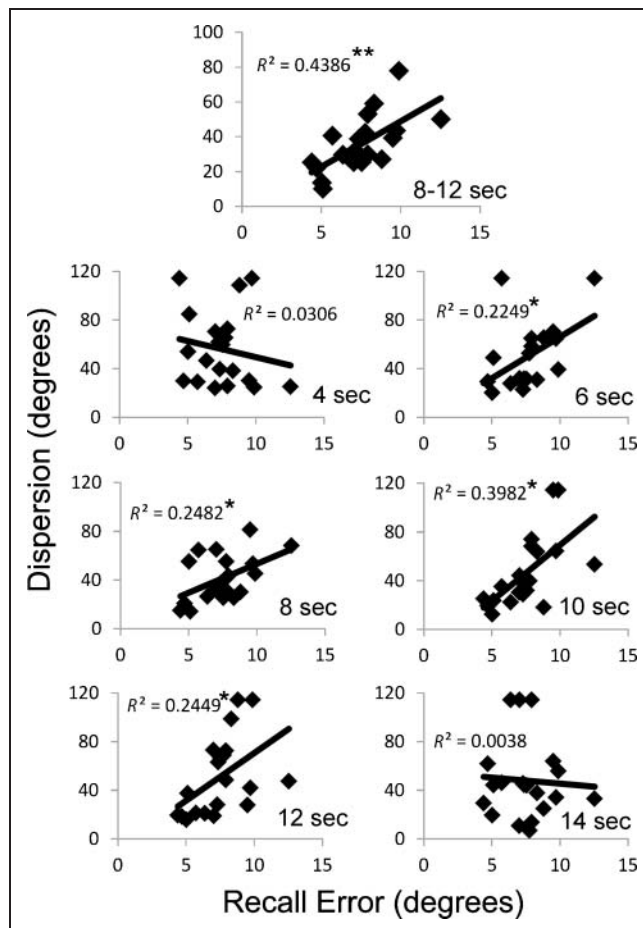


Figure 4. Individual differences in CRF dispersion are strongly correlated with memory performance. Each panel plots mean CRF dispersion (ordinate) from visual areas V1 and V2v as a function of memory performance (mean recall error, abscissa) during different temporal windows (e.g., 8–12 sec; top) or time points (4, 6, 8, 10, 12, or 14 sec; bottom) following the start of each trial. * $p < .05$, ** $p < .01$. Qualitatively similar results were obtained when V1 and V2v were considered separately (see text).




Figure 5. Individual differences in CRF amplitude are uncorrelated with memory performance. Plots mean CRF amplitude (ordinate) from visual areas V1 and V2v (ordinate) as a function of memory performance (mean recall error, abscissa) during different temporal windows (e.g., 8–12 sec; top) or time points (4, 6, 8, 10, 12, or 14 sec; bottom) following the start of each trial. Qualitatively similar results were obtained when V1 and V2v were considered separately.

recall error in every visual area that we examined (Figure 5, top) at any point during the delay interval (Figure 5, bottom).

Experiment 2

One concern is that our findings reflect lingering sensory activity related to stimulus encoding rather than WM storage. To examine this possibility, participants in a second experiment received a cue at the offset of the sample stimulus that instructed them to either store or drop the presented item. To ensure that an adequate number of “store” and “drop” trials were obtained, only five sample orientations were presented (0–144° in 36° increments). Critically, we observed a strong interaction between orientation channel responses and participants' intent to store the sample item, $F(2, 14) = 3.88$, $p = .045$ (Figure 6). When participants voluntarily maintained the sample orientation, channel responses peaked over

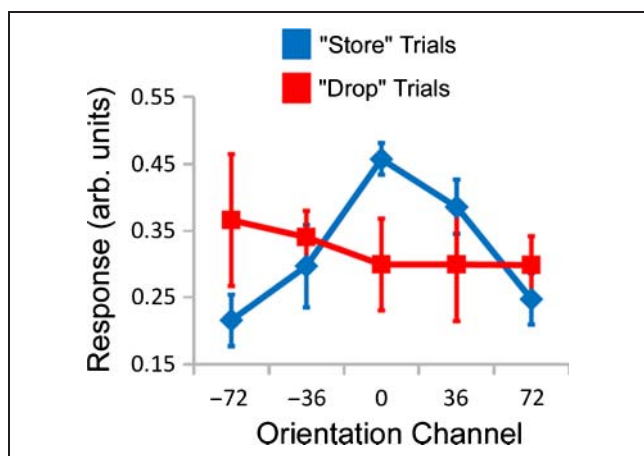


Figure 6. CRFs are contingent on observers' intent to remember the sample orientation. In a second experiment, the sample orientation was followed by a postcue (a change in the color of the fixation point) that instructed observers to remember the sample orientation ("store" trials) or passively wait for the next trial to begin ("drop" trials). A forward encoding model was used to construct CRFs based on patterns of activation observed in V1 and V2v during a period from 8 to 12 sec following the start of each trial. During "store" trials, activation peaked in the orientation channel (0°) and fell off in a graded fashion as distance from this orientation increased. During "drop" trials, this pattern was eliminated. Error bars are ± 1 SEM. Data have been pooled and averaged visual areas V1 and V2v; qualitatively similar results were observed when each area was considered separately.

the stored angle and declined as the distance from this angle increased.² However, when participants received the drop cue, this profile was eliminated. Because the sample period was identical during store and drop trials, these data demonstrate that the results of our first experiment cannot be explained by lingering encoding-related activity.

DISCUSSION

Studies of single unit properties in nonhuman primates have provided the foundation for our understanding of early visual cortex. However, because large populations of neurons are typically involved in the cortical representation of simple stimulus properties, there is great value in developing expedient methods for characterizing population-level tuning functions across broad swaths of the relevant cortical areas. The forward encoding model used here accomplishes this goal and provides a means for quantifying feature-selective tuning functions in human observers engaged in complex behavioral tasks (Brouwer & Heeger, 2011). In this regard, our findings represent a significant advance over the basic observation that sustained patterns of activation in sensory cortices discriminate specific feature values (typically orthogonal directions of motion or orientations, e.g., 45° or 135°) stored in WM (Harrison & Tong, 2009; Serences et al., 2009). Moreover, given that the observed response profiles are linked with mnemonic acuity, this approach can provide a valuable tool for bridg-

ing across human and animal studies to determine how high-acuity representations are maintained in WM.

The current findings demonstrate that intersubject variability in CRF dispersion predicts which individuals will have the greatest memory acuity. However, they do not establish whether moment-to-moment fluctuations in dispersion are linked to memory performance on a within-subject basis. Unfortunately, we were unable to obtain stable estimates of amplitude and dispersion for CRFs observed on individual trials due to large amounts of noise in responses measured on individual trials.³ Thus, whether trial-by-trial fluctuations in dispersion can be used to predict variability in an observer's memory performance remains unclear. Nevertheless, our findings provide critical supporting sensory recruitment models of memory by demonstrating that intersubject differences in the dispersion of sustained stimulus-specific activation patterns observed in visual cortex during WM storage are correlated with variability in mnemonic acuity.

Although the results of Experiment 2 demonstrate that CRFs are contingent on an observer's intent to store information, it is unclear whether the link between CRF dispersion and memory performance described here reflect limiting factors that arise during stimulus encoding, memory storage, or both. By necessity, the amount of information extracted from the stimulus during encoding establishes a "baseline" level of precision that cannot be exceeded (assuming alternative coding strategies have been effectively discouraged). Once this "baseline" representation is established, however, additional factors germane to storage (e.g., decay or interference from other stimuli) could degrade the quality of the representation further. Note, however, that the current experiment was designed to minimize encoding difficulty by presenting the sample stimulus at the fovea for a relatively long interval (1000 msec). Thus, we suspect that the relationship between CRFs and behavior described here reflect the precision of memory storage than encoding per se. However, the current data cannot directly resolve this issue.

Broadly speaking, the fidelity of a population code can be enhanced either by increasing amplitude or reducing dispersion (e.g., Butts & Goldman, 2006). For example, increasing the amplitude should improve the quality of a population code by increasing the signal-to-noise ratio and decreasing the dispersion should decrease the likelihood of reporting an orientation different from the sample. Thus, one lingering question concerns why dispersion, but not amplitude, predicts memory performance. Although we cannot offer a detailed explanation of this finding, we suspect that it might reflect the nature of the observer's task. For example, optimized dispersion might be especially important when observers are required to make very fine-grained discriminations based on information stored in WM (as in the current study). We emphasize that this account is speculative, and further research is needed to explore putative links between dispersion, amplitude, and the precision of memory representations.

Nevertheless, the current study accomplishes the important step of linking mnemonic acuity with a specific characteristic of sustained population-level responses in visual cortex.

Finally, correlations between neural activity and behavioral outcomes cannot provide conclusive evidence that the neural activity plays a causal role in the maintenance of online memory representations. However, establishing a direct link between brain activity and overt behavioral success is a necessary step in developing a neural model of memory that can elucidate the determinants of superior memory performance. Mounting evidence (e.g., Anderson, Vogel, & Awh, 2011; Ester, Vogel, & Awh, 2011; Fukuda, Awh, & Vogel, 2010; Barton, Ester, & Awh, 2009; Zhang & Luck, 2008; Awh, Barton, & Vogel, 2007) suggests that WM ability is determined by two independent factors: the number of items an observer can store, and the precision with which this information can be retained. Significant progress has been made in developing neural measures that are sensitive to individual differences in the number of items that can be simultaneously maintained in visual WM (e.g., Todd & Marois, 2004; Vogel & Machizawa, 2004), but comparatively little is known about the neural mechanisms that mediate the quality of WM representations. Thus, the current findings provide an important complement to past efforts by providing a robust neural measure of individual differences in mnemonic acuity.

Acknowledgments

This study was supported by NIMH R01-MH087214 to E. A. **Contributions:** E. F. E., E. A., and J. T. S. conceived and designed the experiment. E. F. E. and D. A. collected and analyzed the data using software developed by J. T. S. All authors contributed to writing the manuscript.

Reprint requests should be sent to Edward F. Ester, Department of Psychology, University of California, San Diego, 9500 Gilman Drive, Mail Code 0109, La Jolla, CA 92093, or via e-mail: ester@ucsd.edu or Edward Awh, Department of Psychology, University of Oregon, 1227 University of Oregon, Eugene, OR 97403, or via e-mail: awh@uoregon.edu.

Notes

1. We report r^2 values from standard linear regression. However, all critical correlations replicated when r^2 was computed using a robust fitting algorithm (specifically, Matlab's "robustfit" function) that minimizes the influence of prospective outliers.
2. We attempted to fit the observed response profiles with a Gaussian (as in Experiment 1) but were unable to obtain accurate fits for three of eight participants (perhaps because only five unique orientation values were used). Related analyses—including correlating recall error with the slope of the channel response function—were also unsuccessful. Nevertheless, the data are informative insofar as they rule out a purely sensory interpretation of our findings.
3. In one analysis, we divided each observer's neural data into "high" and "low" error bins based on a median split of his or her recall errors over all trials. We then selected a random subset of

20 observers (with replacement), computed the mean CRF within each error bin and fit the resulting profiles, and computed the difference in dispersion within the high- and low-error bins. This procedure was repeated 10,000 times, yielding a distribution difference values. To estimate the distribution of responses obtained under the null hypothesis (i.e., no true difference between dispersion values within the high- and low-error bins), the same procedure was repeated after randomly assigning trials to high- and low-error bins. Direct comparison of the empirical and null distributions revealed a modest trend toward larger dispersion values in the high- relative to low-error bin (one-tailed t test; $p = .08$), but this result varied substantially across different analytical parameters (e.g., the bin widths of the two histograms and the number of voxels included in the modeling).

REFERENCES

- 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.
- Barton, B. B., Ester, E. F., & Awh, E. (2009). Discrete resource allocation in visual working memory. *Journal of Experimental Psychology: Human Perception & Performance*, *35*, 1358–1367.
- Brainard, D. H. (1997). The psychophysics toolbox. *Spatial Vision*, *10*, 433–436.
- Brouwer, G. J., & Heeger, D. J. (2009). Decoding and reconstructing color from responses in human visual cortex. *Journal of Neuroscience*, *29*, 13992–14003.
- Brouwer, G. J., & Heeger, D. J. (2011). Cross-orientation suppression in human visual cortex. *Journal of Neurophysiology*, *106*, 2108–2119.
- Butts, D. A., & Goldman, M. S. (2006). Tuning curves, neuronal variability, and sensory coding. *PLoS Biology*, *4*, e92.
- Cowan, N., Elliott, E. M., Saults, J. S., Morey, C. C., Mattox, S., Hismjatullina, A., et al. (2005). On the capacity of attention: Its estimation and its role in working memory and cognitive aptitudes. *Cognitive Psychology*, *51*, 42–100.
- D'Esposito, M. (2007). From cognitive to neural models of working memory. *Philosophical Transactions of the Royal Society London, Series B, Biological Sciences*, *362*, 761–772.
- Ester, E. F., Serences, J. T., & Awh, E. (2009). Spatially global representations in human primary visual cortex during working memory maintenance. *Journal of Neuroscience*, *29*, 15258–15265.
- Ester, E. F., Vogel, E. K., & Awh, E. (2011). Discrete resource limits in attention and working memory (pp. 99–110). In M. Posner (Ed.), *Cognitive Neuroscience of Attention*. New York: Guilford.
- Fukuda, K., Awh, E., & Vogel, E. K. (2010). Discrete capacity limits in visual working memory. *Current Opinion in Neurobiology*, *20*, 177–182.
- Harrison, S. A., & Tong, F. (2009). Decoding reveals the contents of visual working memory in early visual areas. *Nature*, *458*, 632–635.
- 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*, 5154–5167.
- Pasternak, T., & Greenlee, M. W. (2005). Working memory in primate sensory systems. *Nature Reviews Neuroscience*, *6*, 97–107.
- Pelli, D. G. (1997). The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spatial Vision*, *10*, 437–442.

- Postle, B. R. (2006). Working memory as an emergent property of the mind and brain. *Neuroscience*, *139*, 23–38.
- Riggall, A. C., & Postle, B. R. (2012). The relationship between working memory storage and elevated activity as measure with functional magnetic resonance imaging. *Journal of Neuroscience*, *32*, 12990–12998.
- Ringach, D. L., Shapley, R. M., & Hawken, M. J. (2002). Orientation selectivity in macaque V1: Diversity and laminar independence. *Journal of Neuroscience*, *22*, 5639–5651.
- Serences, J. T., Ester, E. F., Vogel, E. K., & Awh, E. (2009). Stimulus-specific delay activity in human primary visual cortex. *Psychological Science*, *20*, 207–214.
- Serences, J. T., & Saproo, S. (2012). Computational advances towards linking BOLD and behavior. *Neuropsychologia*, *50*, 435–446.
- Super, H., Spekreijse, H., & Lamme, V. A. F. (2001). A neural correlate of working memory in the monkey primary visual cortex. *Science*, *293*, 120–124.
- Todd, J. J., & Marois, R. (2004). Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature*, *428*, 751–754.
- Vogel, E. K., & Machizawa, M. G. (2004). Neural activity predicts individual differences in working memory capacity. *Nature*, *428*, 748–751.
- Xu, Y., & Chun, M. M. (2006). Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature*, *440*, 91–95.
- Zhang, W., & Luck, S. J. (2008). Discrete fixed-resolution representations in visual working memory. *Nature*, *453*, 233–235.