

NEURAL MECHANISMS OF MNEMONIC PRECISION

by

EDWARD F. ESTER

A DISSERTATION

Presented to the Department of Psychology
and the Graduate School of the University of Oregon
in partial fulfillment of the requirements
for the degree of
Doctor of Philosophy

December 2011

DISSERTATION APPROVAL PAGE

Student: Edward F. Ester

Title: Neural Mechanisms of Mnemonic Precision

This dissertation has been accepted and approved in partial fulfillment of the requirements for the Doctor of Philosophy degree in the Department of Psychology by:

Edward Awh	Chairperson and Advisor
Edward Vogel	Member
Nash Unsworth	Member
Terry Takahashi	Outside Member

and

Kimberly Andrews Espy	Vice President for Research & Innovation/Dean of the Graduate School
-----------------------	--

Original approval signatures are on file with the University of Oregon Graduate School.

Degree awarded December 2011

© 2011 Edward F. Ester

DISSERTATION ABSTRACT

Edward F. Ester

Doctor of Philosophy

Department of Psychology

December 2011

Title: Neural Mechanisms of Mnemonic Precision

Working memory (WM) enables the storage of information in a state that can be rapidly accessed and updated. This system is a core component of higher cognitive function – individual differences in WM ability are strongly predictive of general intelligence (IQ) and scholastic achievement (e.g., SAT scores), and WM ability is compromised in many psychiatric (e.g., schizophrenia) and neurological (e.g., Parkinson's) disorders. Thus, there is a strong motivation to understand the basic properties of this system. Recent studies suggest that WM ability is determined by two independent factors: the *number* of items an individual can store and the *precision* with which representations can be maintained. Significant progress has been made in developing neural measures that are sensitive to the number of items stored in WM. For example, electrophysiological and neuroimaging studies have demonstrated that activity in posterior parietal cortex is directly modulated by the number of items stored in WM and reaches a plateau at the same set size where individual memory capacity is exceeded. However, comparably little is known regarding the neural mechanisms that enable the storage of high-fidelity information in WM. This dissertation describes two experiments that evaluate so-called sensory-recruitment models of WM, where the storage of high-fidelity information in WM is mediated by sustained activity in sensory cortices that

encode memoranda. In Chapter II, functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis were used to demonstrate that sustained patterns of activation observed in striate cortex discriminate specific feature attribute(s) (e.g., orientation) that an observer is holding in WM. In Chapter III, I show that these patterns of activation can be observed in regions of visual cortex that are not retinotopically mapped to the spatial location of a remembered stimulus and suggest that this spatially global recruitment of visual cortex enhances memory precision by facilitating robust population coding of the stored information. Together, these results provide strong support for so-called *sensory recruitment* models of WM, where the storage of fine visual details is mediated by sustained activity in sensory cortices that encode information.

This dissertation includes previously published and co-authored material.

CURRICULUM VITAE

NAME OF AUTHOR: Edward F. Ester

GRADUATE AND UNDERGRADUATE SCHOOLS ATTENDED:

University of Oregon, Eugene, OR
Villanova University, Villanova, PA
Nazareth College of Rochester, Rochester, NY

DEGREES AWARDED:

Doctor of Philosophy, Psychology, 2011, University of Oregon
Master of Science, Psychology, 2006, Villanova University
Bachelor of Arts, Psychology, 2004, Nazareth College of Rochester

AREAS OF SPECIAL INTEREST:

Neural mechanisms of visual perception, attention, and memory

PROFESSIONAL EXPERIENCE:

Graduate Teaching Fellow, University of Oregon, 2006-2007; 2009-2011

GRANTS, AWARDS, AND HONORS:

Ingeborg & Byron Ward Outstanding Thesis Award, Villanova University, 2007

Systems Physiology Training Grant, Institute of Neuroscience, University of Oregon, 2007-2010

Student Travel Award, Organization for Human Brain Mapping, 2010

Student Travel Award, Vision Sciences Society, 2011

Marthe E. Smith Memorial Science Scholarship, University of Oregon, 2011

PUBLICATIONS:

Ester EF, Vogel EK, Awh E. (in press) Discrete resource limits in visual attention and working memory. In: *The cognitive neuroscience of visual attention* (Posner MI, ed), New York: Guilford

Umemoto A, Drew T, Ester EF, Awh E (2010) A bilateral advantage for storage in visual working memory. *Cognition* 117:69-79.

Ester EF, Serences JT, Awh E (2009) Spatially global representations in human primary visual cortex during working memory maintenance. *J Neurosci* 29:15258-15265.

Barton BB, Ester EF, Awh E (2009) Discrete resource allocation in visual working memory. *J Exp Psychol Hum Percept Perform* 35:1359-1367.

Serences JT, Ester EF, Vogel EK, Awh E (2009) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207-214

Folk CL, Ester EF, Troemel K (2009) How to keep attention from straying: Get engaged! *Psychon Bull Rev* 16:127-132.

Ester EF, Awh E (2008) The processing locus of interference from salient singleton distractors. *Vis Cogn* 16:166-181.

ACKNOWLEDGMENTS

I would like to thank the members of my committee, Ed Vogel, Nash Unsworth, and Terry Takahashi, for their guidance and advice during the preparation of this dissertation. I would also like to thank John Serences for providing technical and statistical advice. Finally, very special thanks are due to Ed Awh, whose unwavering guidance and support over the course of my graduate career have been instrumental in helping me grow as a scientist and thinker.

To my parents, who have always encouraged me to follow my dreams.

TABLE OF CONTENTS

Chapter	Page
I. INTRODUCTION	1
What Factors Determine Working Memory Ability?	2
Evidence for a Limit in the Resolving Power of Working Memory.....	6
Neural Evidence for a Fixed Capacity Limit in Working Memory	10
How Are Fine Visual Details Stored in Working Memory?.....	13
II. STIMULUS-SPECIFIC DELAY ACTIVITY IN HUMAN PRIMARY VISUAL CORTEX.....	18
Introduction.....	18
Method	20
Observers	20
Behavioral Task	21
Staircase Procedure.....	22
fMRI Data Acquisition and Analysis.....	23
Retinotopic Mapping	23
Voxel Selection.....	24
Multivoxel Pattern Analysis	24
Results.....	25
Behavioral Performance.....	25
Analysis of WM-Related Activation in Visual Cortex	26
Discussion.....	34

Chapter	Page
III. SPATIALLY GLOBAL REPRESENTATIONS IN HUMAN PRIMARY VISUAL CORTEX DURING WORKING MEMORY MAINTENANCE.....	36
Introduction.....	36
Materials and Methods.....	38
Subjects.....	38
Working Memory Task.....	38
Perceptual Monitoring (PM) Task.....	40
Staircase Procedure.....	41
fMRI Data Acquisition and Analysis.....	41
Retinotopic Mapping and Voxel Selection.....	42
Multivoxel Pattern Analysis.....	43
Results.....	44
Multivoxel Pattern Analysis.....	45
Event-related Analysis of WM Data.....	51
Similarities between Spatially Global Signals during Perception and WM Maintenance.....	51
Discussion.....	56
IV. GENERAL CONCLUSIONS.....	62
Sustained Changes in the BOLD Signal as a Measure of WM Storage.....	63
Under What Conditions Does Sensory Recruitment Occur?.....	64
Sensory Recruitment in Non-visual Sensory Modalities.....	65
Conclusions.....	67

Chapter	Page
REFERENCES CITED.....	69

LIST OF FIGURES

Figure	Page
2-1. Working Memory Task	22
2-2. Mean amplitude of the blood-oxygen-level-dependent (BOLD) response in primary visual cortex (V1) across the working memory (WM) delay period and intertrial interval (ITI).....	28
2-3. Feature-selective working memory (WM) modulations revealed by multivoxel pattern analysis.....	30
2-4. Comparison of classification accuracy (collapsed across feature dimensions) in two time bins: 4 through 8 s poststimulus and 4 through 10 s poststimulus	32
3-1. Working Memory Task	39
3-2. Multivoxel pattern analysis of WM data	46
3-3. Event-related analysis of WM data	52
3-4. Global orientation-selective patterns are similar during PM and WM maintenance	55

CHAPTER I

INTRODUCTION

Working memory (WM) refers to a temporary, limited-capacity storage buffer that enables the retention of information in a readily accessible and easily updated state. This system is an integral component of higher cognitive function. For example, inter-individual variations in WM ability are strongly correlated with measures of general intelligence (e.g., Stanford-Binet IQ or Raven's Progressive Matrices; Conway, Cowan, Bunting, Theriault, & Minkoff, 2002; Engle, Tuholski, Laughlin, & Conway, 1999, Fukuda, Vogel, Mayr, & Awh, 2010) and reading ability (e.g., verbal SAT scores; Daneman & Carpenter, 1980). Likewise, WM ability is severely disrupted in a number of psychiatric and neurological disorders, including schizophrenia (e.g., Goldman-Rakic, 1994; Gold, Carpenter, Randolph, Goldberg, & Weinberger, 1997; Gold, Wilk, McMahon, Buchanan, & Luck, 2003), major depression (e.g., Channon, Baker, & Robertson, 1993), Parkinson's (Owen, Iddon, Hodges, Summers, & Robbins, 1997), and Alzheimer's (Baddeley, Bressi, Della Sala, Logie, & Spinnler, 1991). These findings provide a compelling motivation to identify and understand the basic factors that determine WM ability.

To date, converging evidence from multiple studies employing a diverse array of methodologies suggests that WM ability is determined by the confluence of two independent factors: the *number* of representations an individual can store, and the *precision* or fidelity of these representations (for a review, see Fukuda, Vogel, & Awh, 2010). In this dissertation, I will develop and test specific hypotheses regarding the neural mechanisms that enable the storage of high-fidelity stimulus representations in WM. In

the interests of brevity I focus on domain of visual WM. However, there is ample reason to suspect that many of the findings discussed in this chapter and those that follow also apply to other sensory modalities (see Chapter IV for an elaboration of this point).

What Factors Determine Working Memory Ability?

Although WM is critical for virtually all forms of online cognitive processing, converging evidence from multiple sources suggests that it has a capacity limit of only 3-4 items (e.g., Awh, Barton, & Vogel, 2007; Irwin, 1992; Luck & Vogel, 1997; Sperling, 1960). In one notable study, Luck and Vogel (1997) asked subjects to remember “sample” arrays containing a variable number of colored squares over a brief delay. At the end of each trial a test array was presented, and subjects were asked to determine whether the initial or test arrays were identical or differed in the color of a single item. Performance on this so-called “change detection” task was virtually perfect for arrays containing fewer than 3-4 items, but decreased monotonically once this range was exceeded. Based on this profile, Luck and Vogel estimated that subjects could remember, at most, 3-4 items from the sample array. Critically, similar estimates were obtained when the sample array contained more “complex” items such as colored, oriented bars of different sizes. Thus, Luck and Vogel concluded that WM can store about 3-4 integrated objects, regardless of their complexity.

Although it is generally agreed that WM is subject to some form of capacity limit, there is substantial debate regarding how to best characterize this limit. On the one hand, some authors have advocated a “discrete resource” model of capacity, where WM is conceptualized as a limited number of “slots”, each capable of storing a single integrated

object or “chunk” of information (e.g., Barton, Ester, & Awh, 2009; Luck & Vogel, 1997; Rouder et al., 2009; Zhang & Luck, 2008). In contrast, others have advocated a “flexible resource” model of WM capacity, where WM is conceptualized as a pool of mnemonic “resources” that can be allocated to a variable number of items (Bays & Husain, 2008; Bays, Catalao, & Husain, 2009; Bays, Wu, & Husain, 2011; Wilken & Ma, 2004). By this account, there is no fixed limit in the number of items that can be held in WM; individuals can choose to store as few or as many items as they like. However, as more items are stored each receives a smaller proportion of available resources, and memory performance decreases. Conversely, when fewer items are stored each receives a greater proportion of resources, and memory performance is enhanced (relatively speaking).

In an influential study, Wilken and Ma (2004) attempted to distinguish between flexible- and discrete resource models of memory by asking subjects to perform a change detection task similar to the one described by Luck and Vogel (1997). In addition to reporting whether a change did or did not occur, subjects were also asked to indicate their confidence on a four-point scale (from “very confident” to “very unconfident”). This enabled the authors to generate a series of receiver operating characteristics (ROC) that were subsequently fit with quantitative models derived from core prediction of flexible- and discrete resource models of WM capacity. In their formulations, Wilken and Ma characterized the discrete resource model as “high-threshold”, such that an item is either stored perfectly (i.e., with no internal noise) or not at all. Moreover, provided that a probed item was stored in memory, it is assumed that changes will be detected with 100% accuracy. Thus, this model predicts that task performance will be perfect for sub-capacity

arrays (e.g., 1-3 or 4 items), but fall monotonically with increases in set size beyond this range.

Wilken and Ma compared this high-threshold discrete resource model to two different flexible resource models. These models varied in their formulations, but both shared the core assumption of a linear relationship between the number of representations in WM and the internal noise associated with each representation, with no limit in the number of representations that can be simultaneously maintained. The results of this study indicated that the flexible resource model of WM capacity provided an excellent description of the empirically observed data, while the discrete resource model provided a comparably poor fit. Thus, Wilken and Ma concluded that – consistent with the core predictions of the flexible resource model– there is no fixed limit in the number of items that can be stored in WM. Instead, observers can store as many items as necessary, with the caveat that the internal noise associated with each representation held in WM increases as more items are stored.

Other recent studies (e.g., Bays & Husain, 2008; Bays, Catalao, & Husain, 2009; Bays, Wu, & Husain, 2010) have also reported evidence consistent with a flexible resource model of WM capacity. In one example, Bays and Husain (2008) asked subjects to remember the locations or orientations of multiple stimuli (colored squares and oriented bars, respectively) across a short delay. Changes occurred on each trial, and subjects were required to report the direction of change (e.g., leftward or rightward displacement in the location memory task, clockwise or counterclockwise rotation in the orientation memory task). To discriminate between flexible and discrete resource models of WM capacity, Bays and Husain expressed their data as the frequency of a given

response (e.g., leftward displacement or clockwise rotation) as a function of both set size and change size (i.e., the magnitude of probe displacement or rotation relative to the corresponding item in the sample array). These profiles were then fit with cumulative response functions, and set-size dependent changes in the slopes of these functions were then examined with reference to core assumptions of flexible- and discrete resource models of WM. As in the Wilken and Ma study, Bays and Husain (2008) characterized the discrete resource model as high-threshold: an item is stored perfectly, or it is not stored at all. Thus, assuming a capacity limit of 4 items, this model predicts no change in slope from set sizes 1-4, followed by monotonic decreases in slope at larger set sizes. In contrast, a flexible resource model predicts monotonic reductions in slope across all set sizes, because any increase in set size requires subjects to spread mnemonic resources across a greater number of stimuli. The results of this study revealed monotonic reductions in slope across the entire range of tested set sizes (1, 2, 4, or 6) for both tasks (orientation and location), consistent with the predictions of a flexible resource model.

The findings reported by Wilken and Ma (2004) and Bays and Husain (2008) are consistent with the core predictions of a flexible resource model of memory. However, note that both of these studies compared the flexible resource model with a “high-threshold” instantiation of the discrete resource model, where items are assumed to be stored perfectly (i.e., with no internal noise) or not at all. As many others (e.g., Fukuda, Vogel, & Awh, 2010) have noted, this characterization overlooks multiple studies demonstrating that representations held in WM have limited resolution or clarity (e.g., Awh, Barton, & Vogel, 2007; Barton, Ester, & Awh, 2009; Zhang & Luck, 2008). Thus, one can argue the studies reported by Wilken and Ma (2004) and Bays and Husain (2008)

failed to provide a “fair” comparison of flexible and discrete resource models. Below, I review evidence consistent with the hypothesis that representations held in WM have limited resolution or clarity and describe the results of a recent study that compared the flexible resource model of WM capacity with a discrete resource model that acknowledges the limited resolving power of WM.

Evidence for a Limit in the Resolving Power of Working Memory

A number of studies have documented substantial reductions in WM capacity with increasing stimulus complexity (e.g., Alvarez & Cavanagh, 2004; Eng, Chen, & Jiang, 2005). In one notable example, Alvarez and Cavanagh (2004) asked subjects to perform a change detection task that utilized stimuli that varied in complexity (e.g., colored squares, Chinese characters, and three-sided shaded cubes). Complexity was operationalized using an independent visual search task that required subjects to report the presence or absence of a target among distractors from the same stimulus class; stimuli with greater search slopes were deemed more complex. The results of this study revealed a strong inverse relationship between stimulus complexity and change detection performance. For example, Alvarez and Cavanagh reported that subjects could remember approximately 3-4 “simple” stimuli such as colored squares (commensurate with the findings reported by Luck & Vogel, 1997), but only 1-2 “complex” stimuli such as Chinese characters.

In a subsequent study, Awh, Barton, and Vogel (2007) raised the possibility that reductions in change detection accuracy for complex stimuli reflect errors in detecting relatively small changes rather than a reduction in WM capacity per se. To test this

hypothesis, Awh et al. asked subjects to perform a change detection task that utilized arrays containing a mixture of complex stimuli (e.g., Chinese characters and shaded cubes). On 50% of change trials, a Chinese character was replaced with a shaded cube (or vice versa; so-called “cross-category” changes), while on the remaining 50% of change trials, a Chinese character or shaded cube was replaced with a probe from the same stimulus category (“within-category” changes). Awh et al. reasoned that if reductions in change detection performance for complex stimuli are due to errors in detecting relatively small changes, then decreasing the similarity between the sample and test stimuli (via cross-category changes) should have a beneficial effect on change detection performance (relative to performance with within-category changes). In fact, change detection accuracy was relatively poor for within-category changes, replicating the core findings of Alvarez and Cavanagh (2004). However, change detection accuracy for cross-category changes was subsequently equivalent to that observed in a color change detection task similar to the one reported by Luck and Vogel (1997). Thus, when sample-test similarity was minimized, change detection performance was equivalent for simple and complex objects. These findings suggest that the reductions in change detection accuracy with increasing stimulus complexity observed by Alvarez and Cavanagh (2005; see also Eng, Chen, & Jiang, 2005) were due to errors in detecting relative small changes than a reduction in WM storage capacity.

One important implication of the findings reported by Awh et al (2007) is that the change detection task may measure unique aspects of memory ability depending on whether subjects are required to make fine-grained or coarse discriminations between sample and test stimuli. For example, when relatively coarse discriminations are required,

this procedure provides a reliable estimate of the number of items an individual can hold in WM (e.g., Luck & Vogel, 1997; Vogel & Machizawa, 2004). However, when subjects must make fine-grained discriminations between similar sample and test stimuli, then performance may also be limited by the relative clarity or *resolution* of representations held in memory. Critically, these two factors (i.e., the number of items an individual can represent in WM and the relative precision of these representations) appear to be independent aspects of memory ability. For example, Awh et al. (2007) reported that change detection accuracy was reliably correlated across so-called “number-limited” conditions where sample-test similarity was low. Likewise, accuracy was also reliably correlated across “resolution-limited” conditions where sample-test similarity was high. However, performance in number-limited conditions did not correlate with performance in the resolution-limited conditions. In a subsequent study, Fukuda, Vogel, Mayr, and Awh (2010) who asked a large pool of subjects to perform a battery of number- and resolution-limited WM tasks. An exploratory factor analysis on these data identified wholly orthogonal factors for number and resolution, suggesting that number and resolution are indeed distinct aspects of WM ability.

The findings reported by Awh et al. (2007) suggest that WM ability is determined by the confluence of two factors: the *number* of items a person can store, and the clarity or *resolution* of the stored information. This “hybrid” model can be contrasted with a flexible-resource model that assumes an inverse relationship between the number of representations stored in WM and the precision of each representation, with no fixed limit in the number of representations that can be stored. To distinguish between these perspectives, Zhang and Luck (2008) presented subjects with sample arrays of colored

squares followed by a brief retention interval. At the end of each trial, subjects were cued to report the color of a single square by clicking on a color wheel. The logic of this approach was that if the cued item was stored in WM, then subjects should have at least some (i.e., non-zero) information about its color, and their responses should be normally distributed around the true color value. However, if the cued item was not stored in WM, the subject will be forced to guess, leading to a uniform distribution of responses across color space. Using a relatively simple estimation procedure, the density function associated with each type of trial can be recovered. Thus, this procedure enables the simultaneous estimation of (1) the probability that a given item was stored in WM, and (2), the precision with which stored items were represented. According to the hybrid model outlined above, increasing the number of to-be-stored items should decrease the probability that the probed item was stored in memory, but it should only affect the resolution of the stored representations until a relatively small item limit has been exceeded¹. By contrast, the flexible resource model predicts that increasing the number of to-be-stored items should affect the precision of each representation, but there should be no structural limit in the number of items that can receive mnemonic resources. In fact, Zhang and Luck (2008) found that the proportion of subjects' responses attributable to random guessing increased as subjects attempted to store more items in WM. In addition, modest declines in precision were observed until a set size of three items, after which no

¹ This prediction is based on the assumption that subjects can choose to represent the same information in multiple slots. For example, when asked to remember a single oriented bar a subject might choose to store 4 copies of this information (one per slot) and average these representations in an effort to obtain a more precise representation. Likewise, when two oriented bars must be represented the subject might choose to devote two slots to each bar. Here, fewer samples of each stimulus are stored, so resolution will be worse relative to a case where the subject can devote all of their slots to a single item. Thus, this model predicts that estimates of mnemonic resolution will decrease with set size until putative capacity limits are exceeded, after which no further changes in resolution will be observed. See Zhang & Luck, 2008; Anderson et al, 2011 for further discussion of this point.

further declines were observed. Thus, WM storage was constrained by a relatively small item limit, consistent with a hybrid model of WM capacity.

Neural Evidence for a Fixed Capacity Limit in Working Memory

A number of human neuropsychological and neuroimaging studies also suggest that WM has a fixed capacity limit (Anderson et al., 2011; Todd & Marois, 2004; Vogel & Machizawa, 2004; Vogel, McCollough, & Machizawa, 2005; Xu & Chun, 2006). In one example, Vogel and Machizawa (2004) recorded EEG waveforms from subjects as they performed a lateralized WM change detection task that required them to remember a variable number of objects presented in a cued hemifield. These authors observed a sustained, negative voltage wave over posterior contralateral electrodes that began approximately 300 ms following the onset of the memory array and persisted throughout the maintenance period. This so-called contralateral delay activity (CDA) was strongly modulated by the number of objects subjects were required to remember. Specifically, the amplitude of this component increased monotonically from one to three items and reached an asymptotic limit at around four items. Moreover, individual differences in the specific point at which the CDA reached asymptote were strongly predicted by behavioral estimates of WM capacity. Recent neuroimaging studies (e.g., Todd & Marois, 2004; Xu & Chun, 2006) have identified a similar response profile associated localized in human intraparietal sulcus (IPS). For example, Todd and Marois (2004) presented subjects with a variable number of colored discs followed by a short retention interval. Activity in bilateral IPS during this retention interval was strongly modulated by

the number of items subjects were required to remember, reaching an asymptotic limit at around 3-4 items.

Extant neuroimaging data are also broadly consistent with the distinction between number and resolution discussed above. For example, Xu and Chun (2006) asked subjects to remember arrays containing simple or complex shapes and found that activity in inferior IPS scaled with the total number of items subjects were asked remember before reaching an asymptotic limit at around 4 items (consistent with the core finding of Todd and Marois, 2004). Critically, a similar profile was observed when subjects were asked to remember simple or complex stimuli. However, activity in two other cortical areas – the superior IPS and the lateral occipital complex (LOC) – was modulated by the relative complexity of to-be-remembered information. Specifically, when subjects were asked to remember simple objects activity in these regions increased monotonically with set size before reaching an asymptotic limit at around 4 items. In contrast, when subjects were asked to remember complex items activity in these regions reached asymptote at around 2 items. Thus, these findings suggest that independent brain regions track the number and relative complexity of representations stored in WM.

Asymptotic limits in neural activity associated with WM storage are most easily explained by discrete resource models of WM that assume a fixed capacity limit of about 3-4 items. In contrast, it is not immediately clear how a flexible resource model of WM could account for these limits. For example, a core prediction of these models is that mnemonic resources can be used to store as few or as many items as necessary. By this account, subjects should always store all the items in a display, and storage-related neural activity should continue to increase well beyond set sizes of 3-4 items. Furthermore, the

basic finding that CDA and BOLD amplitudes are modulated by the number of items stored in WM poses a challenge to flexible resource models. Specifically, these models predict that the amount of mnemonic resources should always be consumed regardless of how many items the subject is required to store. If CDA amplitude represents the allocation of WM resources – as suggested by strong correlations with individual WM capacity – then CDA amplitude should be equivalent for one-item and three-item arrays. However, this is clearly not the case.

Recent theoretical (Lisman & Idiart, 1995; Raffone & Wolters, 2001) and experimental (Siegel et al., 2009) work has identified a neurally plausible discrete resource model of WM. Specifically, this model assumes each item held in WM is represented though a unique pattern of high-frequency, synchronous firing across large populations of neurons. When multiple items must be held in memory, the high-frequency activity related to each remembered item may be multiplexed within distinct phases of slower oscillatory activity. One attractive aspect of this phase-coding scheme is that it provides a relatively straightforward explanation of discrete capacity limits that have been reported in numerous studies of WM. For example, if information about each item stored in WM must be segregated from the others in a different range of phase orientations, then there should be a maximum number of locations that could be distinctly represented at once. This implies that the discrete resource limit in selection observed in these experiments may ultimately be due to a basic biophysical limitation in how information can be represented in the brain.

How Are Fine Visual Details Stored in Working Memory?

Although the number of items that can be stored in WM is a core limit in human cognition, the evidence reviewed in the preceding section suggests that the quality of representations stored in WM may be best understood as a distinct aspect of memory ability. This distinction is supported by neural evidence suggesting that different brain regions track the number of items stored in WM and the total amount of visual information contained within the stored items (e.g., Xu & Chun, 2006) as well as analyses of individual differences that identify orthogonal factors for number and resolution in WM (Awh et al., 2007; Fukuda et al., 2010). Significant progress has been made in developing neural measures that are sensitive to individual differences in the number of representations that can be simultaneously maintained in WM. For example, multiple electrophysiological (Vogel & Machizawa, 2004) and neuroimaging (Todd & Marois, 2004; Xu & Chun, 2006) studies in humans have demonstrated that activity in regions of posterior parietal cortex increases monotonically with the number of items held in WM before reaching an asymptotic limit at or near behavioral estimates of capacity. However, comparably little is known concerning the neural mechanisms that enable the storage of detailed visual information in WM. Thus, a key goal for future research is to identify and understand the neural mechanisms that determine the clarity or resolution of representations store in WM.

Early investigations concerning the storage of visual information in WM focused primarily on regions of prefrontal cortex (PFC). This focus was driven by a multitude of studies in non-human primates that documented impairments in visual memory following ablation or lesions of PFC (e.g., Goldman & Rosvold, 1970; Gross, 1963; Warren &

Akert, 1964), as well as single-unit recording studies that revealed sustained increases in the firing of PFC cells during the delay period of a delayed-match to sample task (e.g., Fuster & Alexander, 1971; Fuster, 1973; Niki & Watanabe, 1976). However, subsequent studies have made it clear that there is no single brain region or network of brain regions responsible for WM storage. Instead, an emerging view is that WM storage is mediated by the selective and flexible recruitment of cortical areas that have evolved to perform various sensory, cognitive, and motor functions (see Awh & Jonides, 2001; Jonides, Lacey, & Nee, 2005; Postle, 2006; and D'Esposito, 2007 for reviews). For example, single-unit recording studies in macaques have revealed sustained changes in the firing rates of neurons in direction-selective cortical area MT+ when monkeys are required to remember the direction of a motion stimulus over a short delay (e.g., Bisley & Pasternak, 2001), and microstimulation of cortical area MT+ during the delay period of a motion discrimination task has a deleterious effect on memory performance (Bisley, Zaksas, & Pasternak, 2001). Likewise, multiple studies have revealed sustained changes in neural firing rates in object- and face-selective regions of inferotemporal cortex when monkeys are required to remember an object or face over a brief delay (Chelazzi, Miller, Duncan, & Desimone, 1993; Miller, Li, & Desimone, 1993).

A number of human neuroimaging studies have also reported sustained activity in sensory cortices during WM storage. Importantly, the specific sensory cortical regions engaged during WM storage depend strongly on the type of information subjects are asked to remember. For example, a wealth of neuroanatomical and neuropsychological studies suggest that the human (and primate) visual system can be (coarsely) divided into “dorsal” and “ventral” pathways responsible for the processing of spatial and nonspatial

(i.e., object) information, respectively (e.g., Ettlenger, 1990; Ungerleider & Mishkin, 1982). In an early PET study, Courtney, Ungerleider, Keil, and Haxby (1996) asked whether a similar anatomical segregation also exists during the storage of spatial and nonspatial information in WM. In this study, subjects were shown arrays face stimuli that appeared at various spatial locations. In separate blocks, subjects were instructed to remember either the identities of the faces or their positions. Whole-brain PET images were used to measure local changes in regional cerebral blood flow (rCBF) during the delay period of the task. During remember-location blocks, the authors observed sustained increases in rCBF in the superior and inferior parietal cortex, both of which have been implicated in the sensory processing of spatial information. Conversely, during remember-identity blocks, the authors observed sustained increases in a number of occipitotemporal sites, including the fusiform and parahippocampal cortices. Thus, different cortical regions showed evidence of sustained increases in activity depending on what type of information subjects were asked to remember. These findings are consistent with the hypothesis that WM storage is mediated by the selective and flexible recruitment of sensory cortical areas specialized for the processing of specific stimulus attributes.

Psychophysical studies in humans and macaques also suggest that independent, feature-selective mechanisms mediate the storage of elementary visual properties (e.g., color, orientation, spatial frequency) in WM. For example, Mangussen, Greenlee, Asplund, and Dyrnes (1991) asked subjects to remember the spatial frequency of a Gabor over a 10 second delay. A memory mask presented at various intervals during the delay had a deleterious effect on subjects' performance, but only when the spatial frequency of the mask differed from that of the remembered stimulus by an octave or more. Critically,

memory performance was unaffected by the orientation of this mask, consistent with the hypothesis that specialized mechanisms mediate the storage of spatial frequency and orientation information in WM. In another example, Zaksas et al. (2001) trained monkeys to perform a delayed motion discrimination task where the sample and test stimuli either appeared at the same or different locations. These conditions were blocked, so it was always possible to infer the location of the test stimulus given the location of the sample. On some trials, the experimenters presented a random dot mask during the delay period; this stimulus could appear at the location of the sample stimulus, the impending test stimulus, or elsewhere in the display. Task performance was significantly impaired by the presentation of this mask, but only when it appeared at the location of the impending test stimulus. The selective interference produced by this stimulus suggests that memory for the random dot aperture was localized in the cortical regions that were retinotopically mapped to the position of the impending test stimulus. Thus, there appear to be specialized mechanisms that mediate the storage of information regarding stimulus location, orientation, and spatial frequency in WM.

Considered as a whole, the findings discussed in the last few paragraphs motivate the hypothesis that the storage of fine visual details in WM is mediated by *sensory recruitment*, or sustained activity in cortical regions responsible for sensory processing of the memoranda (Awh & Jonides, 2001; Jonides, Lacey, & Nee, 2005; Postle, 2006; and D'Esposito, 2007). In the chapters that follow, I examine two untested predictions of this model. First, the sensory recruitment view assumes that sustained activity in sensory cortices during WM storage reflects the active storage of hypothesis assumes that this activity represents the specific stimulus values of the stored items. Chapter II describes a

study designed to test this hypothesis using fMRI and multivoxel pattern analysis (MVPA). Second, the sensory recruitment view predicts that when required to remember a specific stimulus value such as orientation, humans opportunistically recruit all available sensory machinery in an effort to store this information with maximal precision. Chapter III describes the results of an experiment that provide preliminary insights into this possibility.

Chapter II was previously published with John T. Serences, Edward K. Vogel, and Edward Awh in *Psychological Science*: Serences, JT, Ester, EF, Vogel, EK, & Awh, E (2009) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207-214.

Chapter III was previously published with John T. Serences and Edward Awh in the *Journal of Neuroscience*: Ester EF, Serences JT, Awh E (2009) Spatially global representations in human primary visual cortex during working memory maintenance. *J Neurosci* 29:15258-15265.

CHAPTER II

STIMULUS-SPECIFIC DELAY ACTIVITY IN HUMAN PRIMARY VISUAL CORTEX

This chapter was previously published with John T. Serences, Edward K. Vogel, and Edward Awh in *Psychological Science*: Serences, JT, Ester, EF, Vogel, EK, & Awh, E (2009) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207-214.

INTRODUCTION

Working memory (WM) allows the on-line storage of behaviorally relevant information. One emerging view is that WM is supported by the same neural mechanisms that encode the sensory information being remembered (we term this the sensory recruitment model of WM; see Awh & Jonides, 2001; D’Esposito, 2007; Jonides, Lacey, & Nee, 2005; Postle, 2006). For example, neurons in face-selective regions of inferotemporal cortex show sustained amplitude increases while an observer is holding a face in WM (Chelazzi, Miller, Duncan, & Desimone, 1993; Courtney, Ungerleider, Keil, & Haxby, 1997; Druzgal & D’Esposito, 2001; Lepsien & Nobre, 2007; Miller, Li, & Desimone, 1993; Ranganath, Cohen, Dam, & D’Esposito, 2004). The sensory-recruitment hypothesis assumes that this activity represents the specific stimulus values of the stored items. Here we report a study in which functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA) provided direct support for this

claim, showing that activation patterns in relevant sensory regions represent the specific stimulus value that is held in WM.

MVPA provides a useful tool for identifying the neural regions that mediate WM by focusing on changes in activation patterns as opposed to simply changes in the mean amplitude of the blood-oxygenation-level-dependent (BOLD) response. For example, Offen, Schluppeck, and Heeger (2009) used fMRI to index activation changes in primary visual cortex (V1), a region known to represent orientation and spatial frequency. Although mean response amplitudes in V1 increased during sustained deployments of spatial attention, activation levels were indistinguishable from a low-level baseline when information about orientation (or spatial frequency) was stored in WM. This finding appears to contradict the sensory-recruitment model. However, as Offen et al. noted, neurons that respond preferentially to the remembered orientation should become more active, whereas neurons tuned away from the remembered orientation should be suppressed (relatively speaking; see, e.g., Martinez-Trujillo & Treue, 2004). A differential pattern of activity across the relevant sensory neurons is thought to represent the encoded orientation (Paradiso, 1988; Pouget, Dayan, & Zemel, 2003; Sanger, 1996), and therefore the sensory-recruitment account holds that this pattern should be maintained during a WM delay period as well. However, if the BOLD response spatially integrates information from neurons that are more active (i.e., those tuned to the remembered orientation) with information from neurons that are less active (Logothetis, Pauls, Augath, Trinath, & Oeltermann, 2001), then a failure to find a sustained amplitude increase in orientation-specific regions of cortex during a delay period does not provide strong evidence against the sensory-recruitment model of WM.

We tested the sensory-recruitment hypothesis by determining if WM is mediated by sustained feature-selective activation patterns in cortical regions that process the relevant sensory information. Using fMRI, MVPA, and a pattern-classification algorithm, we examined feature-specific WM modulations in V1 while subjects remembered either an orientation or a color for 10 s (Haxby et al., 2001; Haynes & Rees, 2005; Kamitani & Tong, 2005, 2006; Norman, Polyn, Detre, & Haxby, 2006; Peelen & Downing, 2007; Serences & Boynton, 2007a, 2007b). The observation of feature-specific activation patterns in V1 suggests that sensory mechanisms are recruited to support the storage of information in WM and furthermore indicates that subjects have top-down control over which features of a multifeature stimulus are stored.

METHOD

Observers

Ten neurologically intact observers participated in a single 2-hr scanning session. All observers gave written informed consent in accord with the requirements of the institutional review board at the University of Oregon. Data from 3 observers were discarded because of technical problems or voluntary withdrawal from the study. Each observer was trained in the experimental task for approximately 1.5 hr prior to scanning, to set sample-test disparities to threshold (see Staircase Procedure).

Behavioral Task

Stimuli were rendered on a light-gray background and displayed via a rear-mounted projector (see Figure 2-1). Observers were instructed to maintain fixation on a central square (subtending 1° visual angle from a viewing distance of 58 cm) that was present for the duration of each scan. At the beginning of each trial, observers were shown the sample, a Gabor stimulus (diameter = 13°) with a small circular aperture (diameter = 2°) cut around the fixation square. The sample stimulus was rendered in one of two orientations (45° or 135° , plus or minus an additional offset randomly selected from a range of $\pm 10^\circ$) and in one of two colors (middle red or middle green, plus or minus an offset randomly selected from a range of $\pm 10\%$ saturation). The addition of random jitter around the canonical features discouraged the use of verbal labels, which would not have been precise enough to support accurate performance in this task. For some observers ($n = 2$), the spatial frequency of the stimulus varied randomly across trials (0.75–1.25 cycles/deg); for the others, the spatial frequency was fixed at 1 cycle/deg. The sample stimulus flickered on and off at 5 Hz and was presented for a total of 1 s, followed by a 10-s blank retention interval. Next, observers were shown a test stimulus (again flickering at 5 Hz for a total of 1 s) that either was identical to or mismatched the sample along one feature dimension (i.e., either the orientation or the color was slightly different). Observers then used a custom-made button box to make a two-alternative forced-choice response regarding whether the sample and test stimuli matched. The sample and test stimuli differed on 50% of trials. A 10-s intertrial interval followed the offset of each test stimulus. At the beginning of each block of trials (or scan), observers were informed that the sample and test stimuli would differ along only




Figure 2-1: Working memory task. On each trial, a sample stimulus that flickered at 5 Hz was presented for 1 s; observers were instructed to remember either the exact orientation or the exact color of this sample over the following 10-s delay period. After the delay, a test stimulus was presented, also for 1 s. The task was to indicate with a button press whether or not the test stimulus matched the sample stimulus on the intended dimension. An exaggerated orientation-mismatch trial is depicted here for illustrative purposes. The test stimulus was followed by a 10-s intertrial interval (ITI).

one feature dimension (color or orientation). Thus, remember-orientation and remember-color trials were run in separate scans. Each scan contained eight trials, and each observer completed seven or eight scans in each memory condition. Feedback (percentage correct) was given after each scan.

Staircase Procedure

To encourage a narrow focus of attention on only the relevant feature, and to discourage verbal-labeling strategies, we titrated task difficulty for each observer in a separate behavioral testing session (7–8 blocks of each feature condition). The task was identical to that just described, and the sample-test disparity was adjusted independently for each feature dimension until a criterion level of performance (approximately 75% correct) was reached. The resulting orientation and color disparities were used to determine the sample-test disparity for each feature during scanning.

fMRI Data Acquisition and Analysis

Scanning was performed using a 3-T Siemens Allegra system at the Robert and Beverly Lewis Center for Neuroimaging at the University of Oregon. Anatomical images were acquired using a spoiled-gradient-recalled T1-weighted sequence that yielded images with a 1-mm³ resolution. Whole-brain echo-planar images (EPIs) were acquired in 33 transverse slices (2,000-ms repetition time, 30-ms echo time, 90° flip angle, 64x64 matrix, 192-mm field of view, 3.5-mm slice thickness, no gap). Data analysis was performed using BrainVoyagerQX (Version 1.86; Brain Innovations, Maastricht, The Netherlands) and custom time-series analysis and pattern-classification routines written in MATLAB (Version 7.2; Mathworks, Natick, MA). Either seven (N = 2) or eight (N = 5) scans of the remember orientation and remember-color tasks were collected per subject (14–16 scans per subject), with each scan lasting 186 s. EPI images were slice-time-corrected, motion-corrected (within and between scans), and high-pass-filtered (three cycles per run).

Retinotopic Mapping

Retinotopic maps were obtained using a rotating checkerboard stimulus and standard presentation and analysis techniques (Engel et al., 1994; Sereno et al., 1995). This procedure was used to identify V1, a region known to respond to both color and orientation (e.g., Johnson, Hawken, & Shapley, 2001; Leventhal, Thompson, Liu, Zhou, & Ault, 1995; Sincich & Horton, 2005; Solomon & Lennie, 2007), as well as ventral (V2, V3, hV4) and dorsal (V2, V3, V3a) visual areas.

Voxel Selection

Independent functional localizer scans were used to identify regions of occipital visual cortex that responded to the spatial position occupied by the stimulus aperture in the main experiment. Colored Gabor stimuli identical to those used in the WM task were cycled on and off at 5 Hz for 10 s on each trial and followed by a 10-s passive fixation epoch. At the beginning of each localizer run, observers were instructed to attend to either the color or the orientation of the stimulus for the duration of that run. Color and orientation runs were presented in alternation. During each trial, two or three target events (a change in the value of the relevant feature dimension for 100 ms) occurred, and observers pressed a button whenever they detected a target event. Each localizer run contained 12 trials, and observers completed either one ($N = 4$) or two ($N = 3$) runs per attended feature dimension. We then used a general linear model (GLM) to identify voxels within each visual area that responded more strongly during epochs of stimulation than during epochs of passive fixation. The single regressor in the GLM was created by convolving a boxcar model of the stimulus protocol with a gamma function (Boynton, Engel, Glover, & Heeger, 1996). In each visual area, all voxels that passed a statistical threshold of $p < .05$ (corrected for multiple comparisons using the false discovery rate algorithm in BrainVoyager) were retained for further analysis.

Multivoxel Pattern Analysis

For the MVPA, we first extracted the raw time series from each voxel within each region of interest during a time period extending from 4 s to 10 s after the presentation of each sample stimulus. These time series were then normalized using a z transform on a

scan-by-scan basis. Temporal epochs from all but one scan were extracted to form a “training” data set for the classification analysis; data from the remaining scan were defined as the “test” set (we use the term scan to refer to an entire 186-s data-collection sequence, so the training and test data sets were always independent). We then trained a Support Vector Machine (SVM; specifically, the OSU-SVM implementation, <http://sourceforge.net/projects/svm/>; see also Kamitani & Tong, 2005, 2006) using only the training data and then used the SVM to classify the orientation or the color of the sample stimulus on each trial from the test scan (classification of color and classification of orientation were carried out separately, so chance for all comparisons was 50%).

This procedure was repeated using a hold-one-scan-out cross-validation approach, so that data from every scan were used as a test set in turn. The SVM’s overall classification accuracy for each observer was then defined as the average classification accuracy across all seven or eight permutations of holding one scan out for use as a test set (depending on the number of scans the subject completed). Classification accuracy was averaged across corresponding regions of interest in the left and right hemispheres because no significant differences were observed between left and right visual areas.

RESULTS

Behavioral Performance

Discrimination thresholds did not differ between the two subtypes of either feature dimension (45 vs. 135, red vs. green), and overall accuracy was maintained at the level predetermined by the staircase procedure.

Analysis of WM-Related Activation in Visual Cortex

All analyses reported here are based on the 62 most responsive voxels within each visual area because this was the minimum number of voxels with significant activity across subjects and visual areas; however, our main conclusions are robust even when more or fewer voxels are included in the analysis.

The goal of this study was to use fMRI and MVPA to determine whether, as predicted by the sensory-recruitment hypothesis, there are stimulus-specific modulations in early regions of visual cortex while an observer is remembering a specific orientation or color. Although we examined the response properties of several visual areas in occipital cortex (V1, V2v, V3v, hV4, V2d, V3d, V3a), we focused on V1 because this region contains neurons that are selective for both orientation and color (Johnson et al., 2001; Leventhal et al., 1995; Sincich & Horton, 2005; Solomon & Lennie, 2007), making it the ideal region to test for orientation and color-selective modulations during the storage of information in WM. Independent functional localizer scans were used to identify the 62 V1 voxels in each observer that were most selective for the retinotopic position occupied by the stimulus aperture in the WM task (see the Method section for our voxel-selection logic). Before performing the MVPA, we compared the mean amplitude of the BOLD response in V1 (collapsed across all 62 voxels) during the delay period of the WM task with the mean amplitude of the BOLD response in a corresponding temporal epoch following the test stimulus (see Figures. 2-2a and 2-2b). This comparison controlled for low-level sensory factors, as the sample and test stimuli were essentially identical. Even though response amplitudes were slightly higher 10 s post-sample than 10 s post-test, overall activation levels during the WM delay period

were statistically indistinguishable from activation levels during the corresponding epoch following the test stimulus on both remember-orientation trials and remember-color trials. Two-way repeated measures analyses of variance (ANOVAs) with delay type (WM vs. ITI) and time point (four levels, from 4 s through 10 s) as factors showed no reliable main effect of delay type, $F(1, 6) = 0.4, p = 0.55, \eta^2 = 0.06$, and $F(1, 6) = 0.27, p = 0.62, \eta^2 = 0.04$, respectively. We examined activation 4 s through 10 s post-stimulus because this epoch should reflect activation associated with the retention of information in WM. No other retinotopically organized region showed sustained amplitude increases related to storing information in WM. Because comparing activation during the WM delay period with activation during a “passive” ITI following the presentation of the test stimulus is potentially problematic, we also plotted the mean time series for 24 s following the presentation of each sample stimulus (as in Offen et al., 2009). Activation levels fell back to baseline approximately 8 s after the onset of the sample stimulus (see Figs. 2-2c and 2-2d), in contrast to the sustained amplitude increases often observed in parietal and frontal cortex across the delay period in a WM task (e.g., see Figure 1 in D’Esposito, 2007). Thus, at least in our study, V1 did not exhibit a robust sustained amplitude increase that is often associated with WM maintenance. It is important to note that the lack of amplitude changes in V1 does not rule out the possibility that stimulus-specific patterns of activation are involved in maintaining information in WM. Therefore, we used MVPA to determine if V1 exhibits a stimulus-specific activation pattern during the delay period, consistent with the predictions of the sensory-recruitment hypothesis. The MVPA approach is based on the assumption that some subregions of visual cortex contain submillimeter columns of neurons that are selective for different stimulus

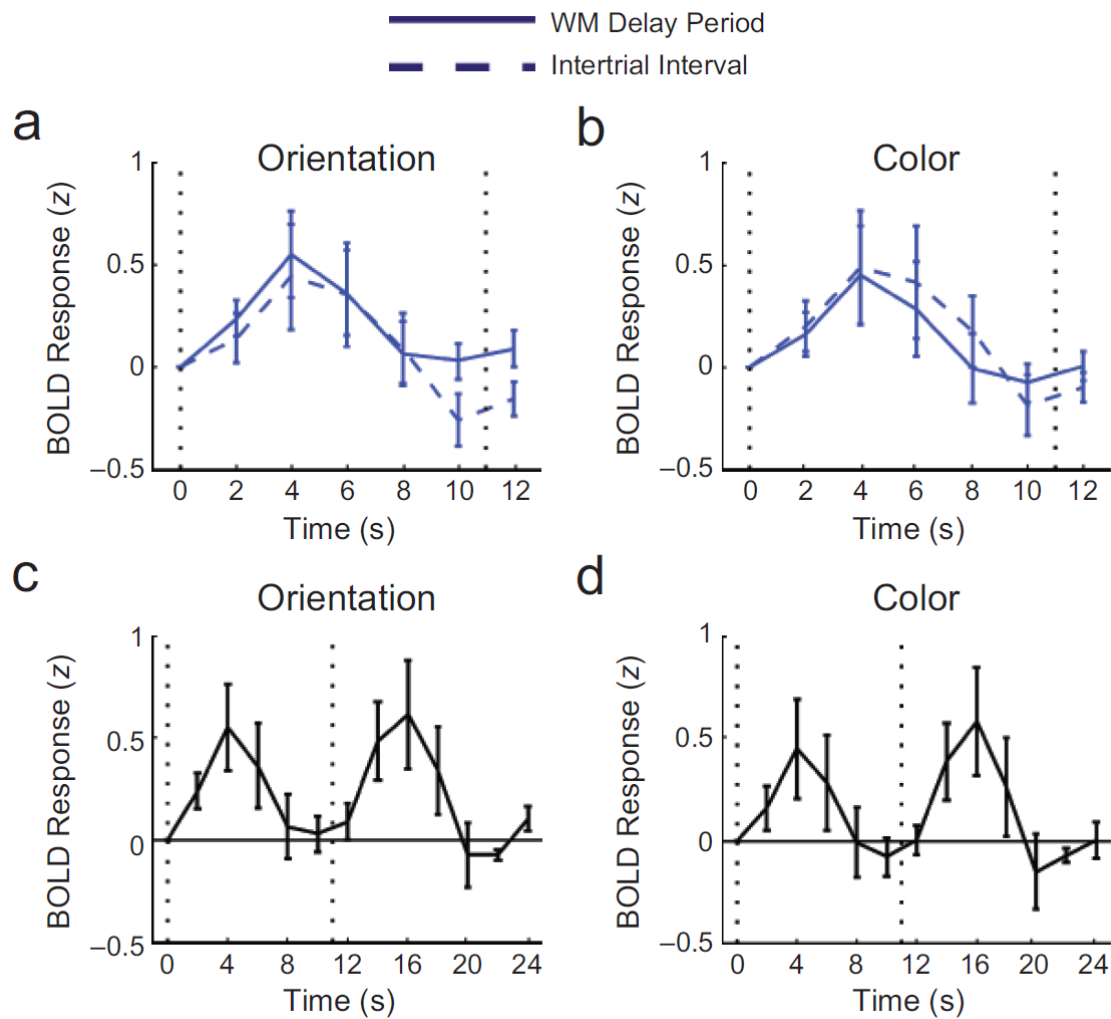


Figure 2-2: Mean amplitude of the blood-oxygenation-level-dependent (BOLD) response in primary visual cortex (V1) across the working memory (WM) delay period and intertrial interval (ITI). All time series were computed against a baseline of the activation level at time 0. The vertical dotted lines highlight the onset of the sample stimulus at 0 s and the onset of the test stimulus at 11 s. The graphs at the top show results for (a) remember-orientation and (b) remember-color trials over a time window extending through 12 s post-stimulus, and the graphs at the bottom show results for (c) remember-orientation and (d) remember-color trials over a longer, 24-s temporal window. Note that because all the event related time series were computed against a baseline of the respective activation level at time 0, the second halves of the time series in (c) and (d) look slightly different from the ITI-evoked responses in (a) and (b) even though they show the same data. Error bars represent ± 1 SEM across observers.

features, such as orientation. In comparison, fMRI voxels are large ($\sim 3 \text{ mm}^3$); however, if slightly more neurons within a voxel prefer a particular orientation than prefer other

orientations, then that voxel may exhibit a weak but detectable response bias (Kamitani & Tong, 2005). By examining the distributed voxel-by-voxel activation pattern across a visual area such as V1, one can make inferences about changes in the underlying population response profile, and pattern-classification algorithms can be used to predict the specific feature that an observer is viewing, attending, or (in our case) remembering (Haynes & Rees, 2005; Kamitani & Tong, 2005, 2006; Norman et al., 2006; Peelen & Downing, 2007; Serences & Boynton, 2007a, 2007b). We therefore examined activation patterns in V1 during the delay period of the WM task (4–10 s following sample onset) to determine if information about the remembered feature was being actively represented. As predicted by the sensory recruitment hypothesis, when observers were remembering the orientation of the sample stimulus, activation patterns in V1 discriminated stimulus orientation, but not stimulus color (see Figure 2-3a). A complementary pattern was observed when observers were instructed to remember the color of the stimulus. A two-way repeated measures ANOVA with memory instruction (remember orientation vs. remember color) and stimulus feature (classify orientation vs. classify color) as factors yielded a significant interaction, $F(1, 6) = 21.4, p < .005, \eta^2 = 0.78$. In contrast, activation patterns associated with a corresponding temporal epoch following the test stimulus—which was physically identical and required a challenging discrimination without storage—did not support above-chance classification accuracy (see Figure 2-3b), $F(1, 6) = 0.14, p = 0.72, \eta^2 = 0.02$, so that there was a three-way interaction of delay type (WM vs. ITI), memory instruction, and stimulus feature, $F(1, 6) = 6.90, p < 0.05, \eta^2 = 0.54$. Given that the test stimulus evoked a BOLD response whose amplitude was statistically indistinguishable from that evoked by the sample stimulus (see Figure 2), these data

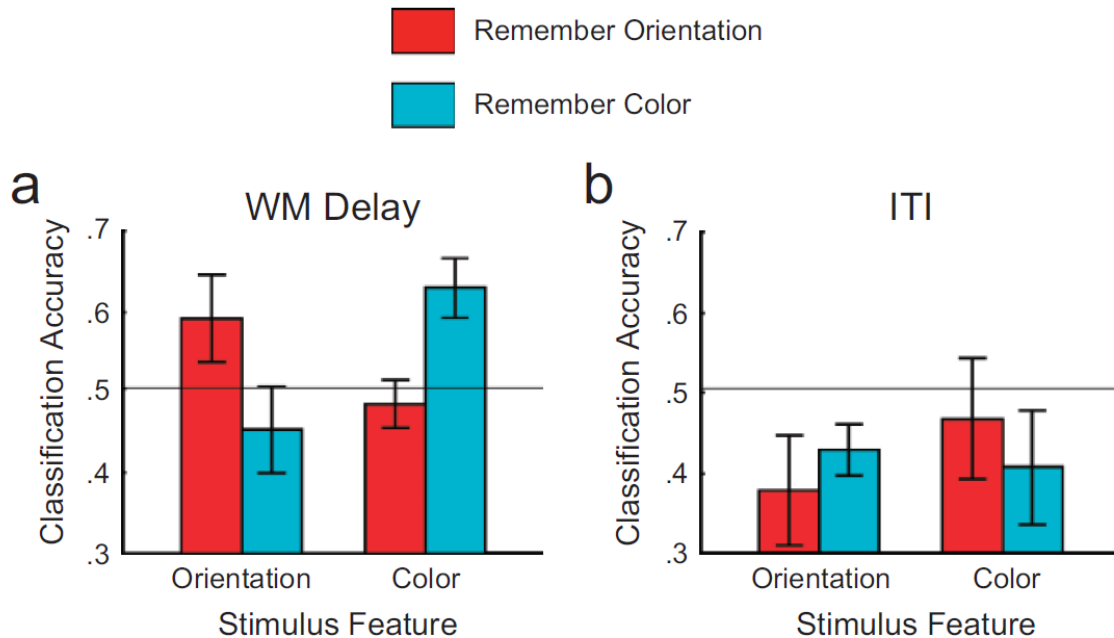


Figure 2-3: Feature-selective working memory (WM) modulations revealed by multivoxel pattern analysis. The graphs show classification accuracy as a function of the stimulus feature (color or orientation) being classified and whether the subject was instructed to remember orientation or color during the scan used as the basis for classification. Results are shown for (a) the WM delay period and (b) an analogous temporal interval following presentation of the test stimulus (i.e., during the intertrial interval, ITI). The horizontal lines at 0.5 accuracy highlight the level of chance performance. Error bars represent ± 1 SEM across observers.

suggest that the active discrimination of the test stimulus for 1 s was not sufficient to drive above-chance classification accuracy. Thus, we conclude that the stimulus-specific pattern of delay activity in V1 was a direct consequence of active maintenance in WM. Activation patterns in other retinotopically organized visual areas did not consistently discriminate the remembered feature of the sample stimulus. The three-way interaction illustrated in Figure 2-3 was also significant when 80 V1 voxels were used to perform the classification, $F(1, 6) = 6.90$, $p < 0.05$, $\eta^2 = 0.53$. Thus, a similar pattern of classification accuracy is observed even when more than 62 voxels are considered in the analysis. Qualitatively similar results were obtained for pattern sizes ranging from 40 to 100

voxels as well. When a linear discriminant classifier based on the Mahalanobis distance between activation patterns was used to compute classification accuracy, the three-way interaction was again significant, $F(1, 6) = 10.80$, $p < 0.025$, $\eta^2 = 0.64$. Thus, the results were not idiosyncratically dependent on the use of an SVM.

At first glance, it is striking that classification accuracy for orientation was below chance when participants were remembering stimulus color (see Figure 3a). However, this effect was not robust across all activation pattern sizes, and we never observed below-chance color classification accuracy when participants were remembering orientation. Together, these analyses suggest that sustained stimulus-specific patterns in V1 reflect active storage in WM, and are not a passive consequence of the attentive encoding of the sample stimulus. Although both the sample and the test stimuli required attentive processing, significant classification accuracy was obtained only during the WM delay period following the sample stimulus. These differences in classification accuracy cannot easily be explained by differences in general arousal or effort related to task demands, as the overall amplitude of the evoked BOLD response was roughly equivalent for the test and the sample stimuli. However, to provide additional support for sustained feature-selective modulations during the WM delay period, we repeated the classification analyses with and without data from the last time point in the delay period (i.e., 10 s after the onset of the sample; see Figures 2-2a and 2-2b). If there were a sustained WM-related activation pattern, then data from the last time point in the delay period would contribute to classification accuracy. By contrast, if feature-selective activation patterns were not sustained across the entire delay period, then adding data from the last time point would not improve classification accuracy (and might even impair classification accuracy if the

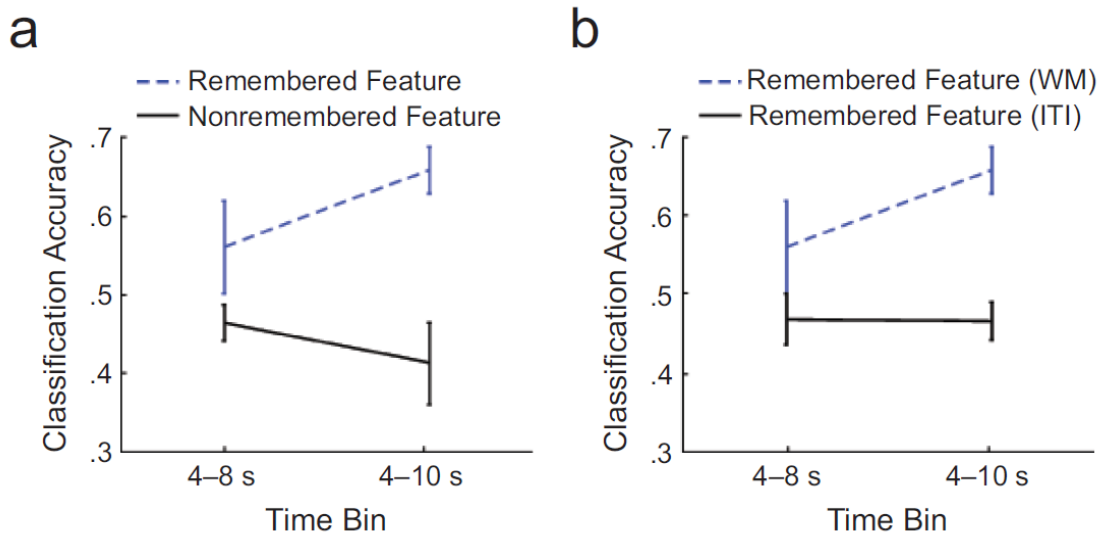


Figure 2-4: Comparison of classification accuracy (collapsed across feature dimensions) in two time bins: 4 through 8 s poststimulus and 4 through 10 s poststimulus. The graph in (a) shows classification accuracy for remembered features (i.e., for orientation when subjects were remembering orientation and for color when subjects were remembering color) and for nonremembered features (e.g., for orientation when subjects were remembering color). The graph in (b) shows classification accuracy for the remembered (or relevant) feature based on data from the working memory (WM) delay period and based on data from the intertrial interval (ITI) following the test stimulus. (Note that following the test stimulus, the “remembered” feature did not need to be remembered anymore and was simply the relevant feature for comparison with the sample.) Error bars represent ± 1 SEM across observers.

activation patterns associated with the last time point were dominated by noise). As Figure 2-4 shows, classification of the remembered stimulus was significantly better when activation patterns associated with the last time point in the delay period were included. However, the addition of this last time point in the delay period did not alter classification accuracy for the non-remembered feature (see Figure 2-4a). Thus, the mere addition of more data did not necessarily improve the performance of the pattern classifier. A two-way ANOVA with memory instruction (orientation vs. color) and time bin (4 s through 8 s vs. 4 s through 10 s) as factors confirmed that adding information from the last time point in the delay period (10 s post-stimulus) selectively enhanced

classification accuracy for the remembered feature, $F(1, 6) = 6.70$, $p < .05$, $\eta^2 = 0.53$. Furthermore, although including data from 10 s post-sample clearly improved classification accuracy for the remembered stimulus feature, including data from 10 s posttest (at the end of the ITI) had little effect on classification of the test stimulus (see Fig. 4b). The two-way interaction of delay type and time bin was significant, $F(1, 6) = 6.20$, $p < .05$, $\eta^2 = 0.51$. Finally, we repeated the analysis after removing data from the peak of the stimulus-evoked BOLD response (4 s; see Figure 2-2) and used only data collected 6 s to 10 s post-stimulus to classify the remembered feature attribute. Even when data from the peak were excluded, classification accuracy for the remembered feature was significantly higher than classification accuracy for the non-remembered feature, collapsed across remember-orientation and remember-color trials (.584 vs. .474), $t(6) = 2.9$, $p_{\text{rep}} > .87$. These control analyses support our conclusion that the interaction depicted in Figure 2-3a reflects the on-line maintenance of information in WM, rather than the aftereffects of a phasic sensory response.

Although our data suggest that maintaining information in WM gives rise to sustained feature-selective activation patterns in V1, a stronger prediction of the sensory-recruitment hypothesis is that the pattern of activation during the delay period will literally mimic the pattern of activation evoked during sensory processing of the same stimulus. To test this prediction, we trained a classification algorithm using data from the functional localizer scans that were initially used to identify visually responsive voxels in each subregion of occipital cortex (see Method). The stimuli used in the localizer scans were identical to those used in the WM study, except that they were presented continuously for 10 s instead of only 1 s, and WM was not required during the localizer

tasks. Each subject completed two to four of these localizer scans (in half of the scans, they attended orientation and ignored color; in the other half, they attended color and ignored orientation). One SVM was trained using data from attend-orientation localizer scans, and another was trained using data from attend-color localizer scans. These SVMs were then used to predict the orientation or color that subjects were remembering on each trial during the main WM task. Collapsed across remember-orientation and remember-color trials, classification accuracy was .607 (SEM = 0.04), $t(6) = 2.50$, $p_{\text{rep}} > .87$. This above-chance classification accuracy demonstrates that the V1 activation pattern that is sustained during WM resembles the sensory-evoked response that is observed during sensory processing alone.

DISCUSSION

These results demonstrate that the maintenance of information in visual WM elicits stimulus-specific activation patterns in the same regions of visual cortex that encode the to-be-remembered sensory information. These activation patterns were specifically tied to the delay period, when active rehearsal in WM was required. Classification accuracy was not above chance following the test stimulus, which involved identical bottom-up stimulation and discrimination of the same stimulus dimension, but no WM load. In addition, the sustained activation patterns observed during the delay period were similar to patterns evoked by the continuous presentation of identical sensory stimuli, which suggests that early feature-selective visual areas are recruited to maintain a “copy” of remembered stimulus attributes, as opposed to a more abstract or categorical

representation. In addition to providing these empirical results, this study demonstrates that MVPA is a valuable tool for answering questions about the neural mechanisms that mediate the storage of specific stimulus values in WM.

Finally, these findings are also relevant to the claim that the capacity of visual WM is determined by the number of individuated objects that have to be stored, rather than the total amount of visual detail contained within those items (Awh, Barton, & Vogel, 2007; Irwin, 1992; Luck & Vogel, 1997; Woodman & Vogel, 2008; Xu & Chun, 2006; Zhang & Luck, 2008). For example, Luck and Vogel (1997) showed that capacity estimates for objects defined by a single feature (e.g., color or orientation) were equivalent to capacity estimates for multifeatured objects (e.g., colored oriented lines). This suggests that capacity is determined by the number of objects that are stored, rather than by the total information load. Alternatively, other researchers have proposed that an obligatory set of core features, including attributes such as color and orientation, is maintained regardless of the observer's intentions (Alvarez & Cavanagh, 2004). This hypothesis might explain Luck and Vogel's observation of equivalent capacity estimates for single and multifeature objects if all the possible features were obligatorily stored even when only a single feature was relevant. However, our results (Fig. 2-3) reveal that early sensory areas selectively represent only behaviorally relevant features during a WM delay period and suggest that observers have top-down control over which features are stored (see also Olivers, Meijer, & Theeuwes, 2006; Woodman & Vogel, 2008).

CHAPTER III

SPATIALLY GLOBAL REPRESENTATIONS IN HUMAN PRIMARY VISUAL CORTEX DURING WORKING MEMORY MAINTENANCE

This chapter was previously published with John T. Serences and Edward Awh in the *Journal of Neuroscience*: Ester EF, Serences JT, Awh E (2009) Spatially global representations in human primary visual cortex during working memory maintenance. *J Neurosci* 29:15258-15265.

INTRODUCTION

Working memory (WM) supports the maintenance of information in an online state. Human neuroimaging (Harrison and Tong, 2009; Serences et al., 2009a) and single-unit studies in non-human primates (Miller et al., 1996; Super et al., 2001) suggest that WM storage is mediated by sensory recruitment, or sustained activity in early cortical regions that encode information. For example, Serences et al. (2009a) required subjects to remember the orientation or color of a foveally presented grating. Using functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA), they demonstrated that sustained patterns of activation in primary visual cortex (V1) during WM storage discriminated the specific value of the remembered attribute, even in the absence of changes in overall response amplitude.

One important question concerns the spatial extent of sensory recruitment. On one hand, recruitment may be confined to cortical regions that were activated during the encoding of a stimulus. Dill and Fahle (1998) showed subjects two dot patterns separated

by a short interval and asked them to determine whether these stimuli were identical. Relative to a condition in which the stimuli appeared in the same spatial location, performance was reliably reduced when they appeared in different locations (see also Zaksas et al, 2001; Hollingworth, 2006; 2007). These findings suggest that object representations held in WM are spatially-specific, thus motivating the hypothesis that sensory recruitment is confined to cortical regions that are retinotopically mapped to the spatial location occupied by the memoranda.

Alternatively, sensory recruitment may extend to cortical regions that were not engaged during stimulus encoding. For example, a specific orientation could be represented via the broad recruitment of cortical neurons that encode orientation, regardless of the retinotopic position of the remembered stimulus. This possibility is suggested by studies demonstrating a spatially global spread of feature-based attention during perception (Treue & Martinez-Trujillo, 1999; Serences & Boynton, 2007a). For example, Serences and Boynton (2007a) instructed subjects to monitor one of two directions of motion carried by stimuli located in one hemifield of a display. Using fMRI and MVPA, these authors demonstrated that patterns of activity in visual areas ipsilateral to the stimuli discriminated the attended direction of motion, despite the fact that these regions were not being driven by a bottom-up signal. During WM maintenance, this kind of spatially global recruitment might enhance the precision of mnemonic representations by recruiting additional neurons that support robust population coding of the stored information (Pouget et al., 2003).

Here, we attempted to determine whether sensory recruitment during WM maintenance is spatially local or global. Subjects were required to remember the

orientation of a grating presented in the left or right visual field. Using fMRI and MVPA, we found that patterns of activity in early cortical regions (V1) both contralateral and ipsilateral to a stimulus discriminated the remembered orientation. Under identical sensory conditions where WM maintenance was not required, these effects were abolished. Furthermore, patterns of activity in contralateral and ipsilateral ROIs were qualitatively similar during perception and WM maintenance, consistent with sensory recruitment accounts of storage in visual working memory.

MATERIALS AND METHODS

Subjects

Twenty neurologically intact subjects participated in a single two-hour scanning session. All subjects gave informed written consent in accordance with the Office for the Protection of Human Subjects at the University of Oregon. Data from three subjects were discarded due to technical difficulties ($n = 2$) or voluntary withdrawal from the study ($n = 1$). Prior to scanning, each subject participated in a 1.5h training session to identify performance thresholds in each behavioral task (see *Staircase Procedure*, below). Subjects were compensated at a rate of \$8 per hour for behavioral testing and \$25 per hour for scanning.

Working Memory Task

Stimuli were generated in Matlab (Version 7.1; Mathworks, Natick, MA) using Psychophysics Toolbox software (Brainard, 1997; Pelli, 1997) and rendered in black on a

medium-grey background via a rear projection system. Subjects were instructed to maintain fixation on a central dot (subtending approximately 0.2° from a viewing distance of 58cm) throughout each scan. Eye position was monitored via an ASL (Bedford, MA) model 5000 eyetracking system.

The sequence of events on each trial is presented in Figure 3-1. At the start of each trial, a “sample” grating stimulus (radius 4° , 1 cycle/deg) was displayed in the upper left or right portion of the display (horizontal and vertical eccentricity of $\pm 8.5^\circ$ and $+5^\circ$, respectively). The sample was rendered with an orientation of 45° or 135° (jittered on a trial-by-trial basis by a randomly selected value in the range of $\pm 10^\circ$ to discourage the use of categorical labels), and cycled on and off at a rate of 2Hz (250ms on, 250ms off) for a total of 6 seconds. To attenuate the potency of retinal afterimages and any perception of

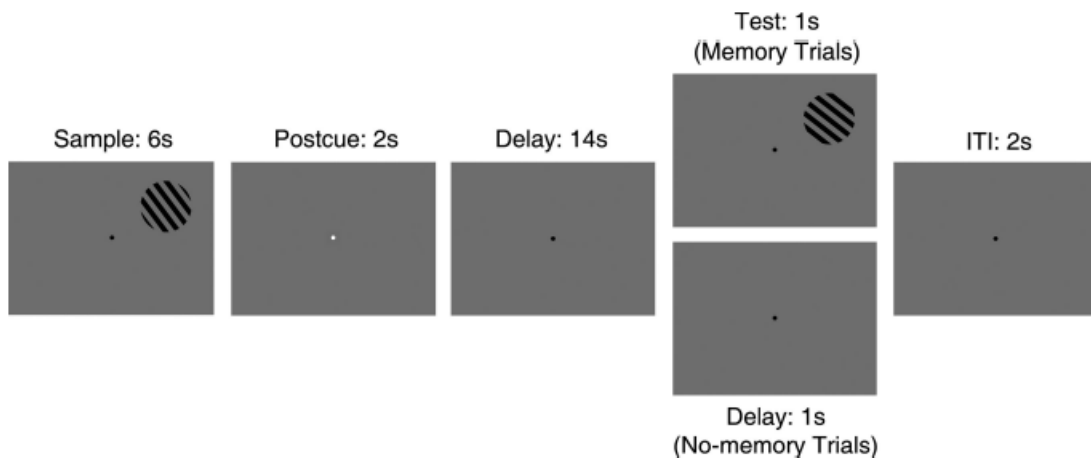


Figure 3-1: Working Memory Task. On each trial, observers were shown a sample grating in the upper left or right visual field for 6 s. A subsequent change in the color of the fixation point (to red or green; shown here in white) informed observers whether they should remember the sample orientation (memory trials) or simply wait for the next trial to begin (no-memory trials). On memory trials, the cue was followed by a 14 s delay period and the presentation of a 1 s test grating in the same spatial location as the sample. Observers made a two-alternative, forced-choice response indicating whether the sample and test gratings shared the same orientation. On no-memory trials, the 14 s delay period was followed by an additional 1 s blank interval (i.e., no test grating was shown and no response was required). Trials were separated by a 2 s intertrial interval (ITI).

apparent motion, the spatial phase of the sample was randomized on each cycle. After the sample epoch, a 2 second change in the color of the fixation point (to red or green) instructed subjects to remember the orientation of the sample for comparison with a subsequent test grating (a “memory” trial) or wait for the start of the next trial (a “no-memory” trial). For nine subjects, a green cue was used to denote memory trials and a red cue denoted no-memory trials. For the remaining eight subjects, this mapping was reversed. The cue was followed by a 14 second delay period. On memory trials, the delay period was followed by the presentation of a static, 1 second test grating and a 2 second inter-trial interval. Subjects made a button press response during this 3-second interval to indicate whether the sample and test stimuli shared the same orientation (50% of trials). On no-memory trials, the delay period was instead followed by an additional 3 seconds of fixation. Stimulus orientation (45° or 135°), position (left or right visual field), and trial type (memory or no-memory) were balanced within each scan (where “scan” refers to a continuous, 404s-long block of 16 trials). Subjects completed a total of 6 (n = 1), 7 (n = 7), 8 (n = 7), or 9 (n = 2) scans as time permitted.

Perceptual Monitoring (PM) Task

During the same scanning session, subjects also performed a perceptual monitoring task. The stimuli and display parameters were closely modeled after those used in the WM task. On each trial, a sample grating oriented to 45° or 135° (randomly jittered on a trial-by-trial basis by a value in the range of $\pm 10^\circ$) was presented in the upper left or right visual field. This stimulus cycled on and off at a rate of 2 Hz for the duration of each 15 second trial. Subjects were instructed to monitor this stimulus and make a

manual button press response whenever they detected a brief (one cycle) change in its orientation (a target event; these occurred at unpredictable intervals 2-3 times per trial; see *Staircase Procedure* for information about the size of these angular deviations). Each 260s-long scan contained a total of 12 real and 3 “null” (i.e., 15 seconds of fixation) trials, and each observer completed a total of 3 ($n = 2$) or 4 ($n = 15$) scans as time permitted.

Staircase Procedure

To ensure that both behavioral tasks were sufficiently challenging, we adjusted their difficulty for each subject during a separate behavioral testing session (completed 1 to 3 days prior to the scanning session). Sample-test (for the WM task) and sample-target (for the PM task) disparities were independently adjusted for each orientation category (45° or 135°) until a criterion level of performance (75% accuracy) was reached. The resulting orientation disparities were used to set sample-test and sample-target disparities during the scan session.

fMRI Data Acquisition and Analysis

fMRI data were collected using a 3T Siemens Allegra system at the Robert and Beverly Lewis Center for Neuroimaging at the University of Oregon. Anatomical images were acquired using a spoiled-gradient-recalled T1-weighted sequence that yielded images with a 1-mm^3 resolution. Whole-brain echo-planar images (EPIs) were acquired in 33 transverse slices (3mm in-plane resolution, 2000-ms TR, 30-ms TE, 90° flip angle, 64×64 matrix, 192-mm FoV, 3.5-mm slice thickness, no gap). EPIs were slice-time

corrected, motion corrected (both within and between scans), and high-pass filtered (three cycles per run). Image preprocessing and data analysis were performed using BrainVoyagerQX (Version 1.9) and custom time-series and pattern-classification routines written in Matlab.

Retinotopic Mapping and Voxel Selection

Retinotopic mapping data were acquired using a rotating checkerboard stimulus flickering at 8 Hz and subtending 45° of polar angle (following Engel et al., 1994; Sereno et al., 1995). Each observer completed one scan lasting 480s. This procedure was used to identify visual areas V1, V2v, V3v, V4v, V2d, V3d, and V3a in each cortical hemisphere. To aid in the visualization of these regions, data were projected onto a computationally inflated representation of each observer's grey/white matter boundary.

To identify spatially selective voxels in these visual areas, we constructed a general linear model (GLM) with a single boxcar regressor (denoting stimulus location, i.e., left vs. right visual field) using data from all four PM scans. This regressor was then convolved with a gamma function to account for the assumed shape of the hemodynamic response (see Boynton, Engel, Glover, & Heeger, 1996). Voxels that showed a stronger response during epochs of contralateral (relative to ipsilateral) stimulation ($p < .05$, Bonferroni corrected) were used to define ROIs in V1, V2v, V3v, hV4v, and V3a (as stimuli were presented in the upper visual field, ROIs in V2d and V3d typically contained too few voxels to enable a meaningful analysis).

Multivoxel Pattern Analysis

MVPA provides a powerful complement to conventional univariate analyses of fMRI data. Posterior cortical regions such as V1 contain submillimeter columns of neurons that are selective for different stimulus features such as orientation. MVPA assumes that if a particular fMRI voxel contains slight preponderance of columns tuned to a specific feature value, it should give rise to a weak but detectable response bias. By considering patterns across multiple weakly selective voxels, it is possible to infer the specific feature values that subjects are attending (e.g., Haxby et al., 2001; Haynes & Rees, 2005; Kamitani & Tong, 2005; Norman et al., 2006; Peelen & Downing, 2007; Serences & Boynton, 2007a; 2007b; Serences et al. 2009b), imagining (Stokes et al., 2009) or remembering (Serences et al., 2009a; Harrison & Tong, 2009), even in the absence of sustained changes in overall response amplitude (Serences et al., 2009a).

To perform MVPA, we first extracted the raw time series from each voxel in a given ROI during a time period extending from 6 to 16 seconds following the offset of the sample stimulus (i.e., 12-22 seconds after the start of each trial). Each time series was normalized on a scan-by-scan basis using a z transform and sorted into one of eight bins according to three factors: ROI (contralateral or ipsilateral to the stimulus location), orientation (45° or 135°), and memory condition (memory or no-memory). We then defined a “training” data set using the data from all but one scan (here, “scan” refers to a single, continuous block of trials; thus, the training and test sets were always independent). Within the training data set, we computed activation patterns comprising the mean response of each voxel during 45° and 135° trials. This was done independently for each location x memory condition pair, yielding a total of 8 activation patterns. Data

from the remaining scan were used to define a “test” data set using an analogous approach. Before classification, we removed the overall mean from each activation pattern in the training and test sets. Finally, we trained a Support Vector Machine (SVM; specifically, the OSU-SVM implementation, <http://sourceforge.net/projects/svm>) to discriminate sample orientation using patterns from the training set, then used it to infer the orientation of the sample stimulus on each trial in the test set. Classification was performed separately for each stimulus location and memory condition pairing, so chance performance was always 50%. This analysis was iterated using a hold-one-out cross-validation procedure (see, e.g., Kamitani & Tong, 2005) until data from every scan had been used as the test set. Overall classification accuracy for each observer was then defined as the average classification accuracy for each observer across all 6, 7, 8, or 9 permutations (depending on how many WM scans the observer completed) of the hold-one-out procedure.

RESULTS

The primary goal of this study was to use fMRI and MVPA to examine the spatial extent of sensory recruitment during WM maintenance. We examined response profiles within several regions of occipital cortex (V1, V2v, V3v, hV4v, V3a); however, we focus on V1 as previous demonstrations show that this region displays robust patterns of orientation-selective activity during both perception (Haynes & Rees, 2005; Kamitani & Tong, 2005; 2006; Serences & Boynton, 2007a; Serences et al. 2009b) and WM maintenance (e.g., Serences et al., 2009a; Harrison & Tong, 2009). For all analyses, data

are presented as collapsed across corresponding ROIs in each visual area (e.g., left and right V1) as no differences in classification accuracy were observed across cortical hemispheres. All analyses reported here are based on the 60 most spatially selective voxels in each ROI. Specific V1 ROIs in three subjects contained fewer than 60 voxels. For these subjects, analyses were performed using all available voxels within each ROI. Critically, we observed no systematic relationship between ROI sizes and overall classification accuracy, indicating that the findings reported here were not unfairly influenced by subjects with larger ROIs. In addition, all findings are robust across a broad range of pattern sizes (i.e., 50 – 75 voxels).

Overall, orientation discrimination thresholds were higher during perceptual monitoring (two-sample paired t-test, $p < .001$), indicating that this task was more challenging than the WM task. However, no differences in discrimination thresholds were observed between orientation categories (45° and 135°) in either task, and overall accuracy remained at a level near that dictated by the staircase procedure.

Multivoxel Pattern Analysis

If sensory recruitment operates in a spatially global fashion, then sustained patterns of activation in ROIs both contralateral and ipsilateral to a remembered stimulus should discriminate its orientation. Figure 3-2 depicts the results of an MVPA analysis based on the 60 most spatially selective voxels (as identified using data from the perceptual monitoring task, see *Retinotopic Mapping and Voxel Selection*, Methods) from each V1 ROI during a time period extending from 6 to 16 seconds following the offset of the sample stimulus. An analysis of variance revealed a main effect of memory condition




Figure 3-2. Multivoxel pattern analysis of WM data. The results of a multivoxel pattern analysis using the 60 most spatially selective voxels in each V1 ROI during a time period extending from 6 to 16 s after the offset of the sample stimulus. The horizontal line at 0.5 denotes chance classification accuracy. ROIs both contralateral and ipsilateral to the sample stimulus discriminated its orientation but only when active WM maintenance was required. Qualitatively similar results were obtained across a range of pattern sizes (50 and 75 voxels) as well as when a linear discriminant algorithm was used to perform classification. Error bars represent ± 1 SEM.

($F_{(1, 16)} = 11.07, p < .001$), no effect of stimulus location (contralateral or ipsilateral; $F_{(1, 16)} = 1.49, p > .24$), and no interaction between these two factors ($F_{(1, 16)} < 1$). Planned comparisons revealed that classification accuracy was higher on memory trials relative to no memory trials. Critically, this was true in both contralateral ($t_{(16)} = 3.48, p < .01$) and ipsilateral ($t_{(16)} = 2.84, p < .05$) ROIs. Moreover, we failed to observe consistent above-chance classification in either ROI on no memory trials ($t_{(16)} = .19, p > .8$ and $t_{(16)} = -1.99, p > .06$ for contralateral and ipsilateral ROIs, respectively; apparent below chance classification in ipsilateral ROIs during no-memory trials was not robust across variations

in pattern sizes). As sensory conditions were equivalent during memory and no memory trials, these findings suggest that the above-chance classification observed in both ROIs during memory trials reflects active WM maintenance rather than a lingering sensory response to the sample stimulus. Qualitatively similar findings were observed across a range of pattern sizes (50, 75 voxels) and classification algorithms (e.g., linear discriminant function), indicating that above-chance classification is not idiosyncratically dependent on the use of a SVM.

MVPA was also applied to data from extrastriate visual areas V2v, V3v, hV4v, and V3a. However, contralateral and ipsilateral ROIs in each of these areas failed to support above chance classification accuracy during memory or no-memory trials (see Serences et al., 2009a, for similar findings using foveal stimuli). Qualitatively similar results were observed over a wide range of pattern sizes (i.e., 50, 60, and 75 voxels) and different classification algorithms (e.g., linear discriminant function). We suspect that this lack of generalization is due to the relatively low salience of our stimuli. For example, our small, peripheral stimuli produced a relatively weak response in only a few dozen voxels within each extrastriate visual area we examined. It is possible that with a larger, sustained response, patterns of activity in each of these visual areas would support robust classification of the remembered orientation. Relevant to this issue, MVPA was also applied to data from the PM task (using the same hold-one-out cross validation procedure described above; see *Multivoxel Pattern Analysis*, Methods). Here, patterns of activity contralateral and ipsilateral V1 ROIs discriminated the orientation of a monitored stimulus. In addition, patterns of activity in several extrastriate visual areas also discriminated the monitored orientation. However, these effects were confined to

contralateral ROIs. We speculate that the absence of above chance ipsilateral classification in extrastriate visual areas is due to the fact that this analysis was based on a small amount of data (each subject completed only 3-4 PM scans).

Our WM findings are consistent with the hypothesis that sensory recruitment during WM maintenance is spatially global. However, there are a number of possible issues that might also account for these effects. One such issue is head motion. For example, subjects may have made small head movements consistent with the remembered orientation (e.g., rightward for 45° trials and leftward for 135° trials). Such movement could differentially alter the intensity of fMRI voxels and bias the results of MVPA. To quantify any differences in head position as a function of the remembered orientation, we compared the mean change (relative to the start of each trial) in three parameters corresponding to leftward and rightward head movement (translation in the axial plane, as well as rotation in the coronal and axial planes) during a window extending from 6-16s following the offset of the sample stimulus (12-22s; the same temporal window used to perform MVPA). Across all three parameters, we failed to observe any differences in head position on 45° and 135° trials (two-sample, repeated measures t-tests; all $ps > .1$). Thus, it is unlikely that the spatially global recruitment effects reported here are due solely to differences in head position.

A second potential issue is eye position. Although subjects were instructed to maintain fixation for the duration of each WM scan (as assessed by visually monitoring ongoing eye position in all subjects with the aid of an ASL tracking system) it is possible that they made subtle eye movements consistent with the orientation of a remembered grating. To further assess compliance with fixation instructions and examine whether

there were any subtle differences in eye position as a function of the remembered orientation, eye position data were recorded as for four subjects during scanning. We computed mean eye position during a period 6-16 seconds following the offset of the sample stimulus on memory trials (the same temporal window used for MVPA). We observed no systematic differences in mean eye position as a function of stimulus orientation (two-sample, repeated measures t-tests, all $ps > .3$), suggesting that the spatially global recruitment effects reported here are not due to different patterns of eye movements on 45° and 135° trials.

Recent human neuroimaging (Serences & Boynton, 2007a) and single-unit (Martinez-Trujillo & Treue, 2004) studies in nonhuman primates have reported spatially global feature-based attention effects during perceptual processing. To examine whether similar effects could be observed in the present study, we performed a classification analysis on patterns observed in each ROI during a period 4-10s following the onset of the sample stimulus. Because the sample epoch of memory and no-memory trials were identical, both types of trial were included in this analysis. In contrast to the spatially global effects we observe during WM maintenance, neither contralateral nor ipsilateral V1 ROIs reliably discriminated the orientation of the sample stimulus (two-tailed t-tests against chance, all $ps > .1$). We suspect that this lack of generality is due to the relatively low salience of our stimuli. For example, MVPA was also applied to data from the PM task using a longer analysis window (6-16s following stimulus onset). Following Serences and Boynton (2007a), patterns of activity in contralateral and ipsilateral V1 ROIs reliably discriminated the orientation of a monitored grating. However, when the same analysis was repeated using a shorter analysis window (4-10s following stimulus

onset), these effects were abolished. We speculate that with a longer sample epoch, patterns of activity in both contralateral and ipsilateral ROIs would support robust classification of sample orientation.

Because classification was carried out independently for each ROI, these findings do not establish whether contralateral and ipsilateral patterns within a given region are similar. To evaluate this possibility, we repeated our classification analysis using data from contralateral memory trials as the training set and examined whether this would allow accurate classification of the stored orientation when the same cortical ROI was ipsilateral to the sample stimulus (using the same hold-one-out cross validation procedure described above; see *Multivoxel Pattern Analysis*, Methods). Across a range of pattern sizes (50, 60, 75 voxels), this analysis failed to reveal above chance classification accuracy (one-sample t-tests against chance, all $ps > .1$). Thus, although patterns of activation in both contralateral and ipsilateral ROIs discriminate the orientation of a remembered stimulus (see Figure 3-2), the pattern of activity in a given ROI is not necessarily identical across trials in which that region is contralateral and ipsilateral to the stored stimulus (see Serences & Boynton, 2007a, for a qualitatively similar finding in the context of a perceptual discrimination task). We speculate that this lack of generalization may be due to differences in lingering bottom-up activity in when a given ROI is contralateral and ipsilateral to the stored stimulus. However, it is also possible that patterns of activation in ipsilateral ROIs may be qualitatively different from those observed in contralateral ROIs.

Event-related Analysis of WM Data

Data from the WM task were also submitted to a conventional univariate analysis. Figure 3-3 depicts the mean amplitude of the BOLD response in each V1 ROI during both memory and no-memory trials (no significant differences in mean response amplitude as a function of angle were observed; thus the data are shown collapsed across this factor). To quantify differences between memory conditions and ROIs during WM maintenance, we computed the mean BOLD response within each condition during a period extending from 6-16s following the offset of the sample stimulus (12-22s in Figure 3-3). A 2 (memory condition) x 2 (ROI) ANOVA of these data revealed a main effect of memory condition, $F_{(1, 16)} = 11.98, p < .01$, no effect of ROI, $F_{(1, 16)} < 1$, and no interaction between these factors, $F_{(1, 16)} = 3.65, p = .07$. Curiously, mean response amplitudes on no-memory trials were higher than those on memory trials ($M = .15$ vs. $-.15$, respectively; standardized units), particularly towards the end of the memory period. Qualitatively similar effects were also observed in extrastriate visual areas.

Similarities between Spatially Global Signals during Perception and WM Maintenance

Recent studies (Serences et al., 2009a; Harrison & Tong, 2009) have reported a high degree of similarity between stimulus-specific patterns of activation in early sensory regions during perceptual processing and WM maintenance, consistent with the hypothesis that WM involves the ingeneration of a specific perceptual event. For example, Serences et al. (2009a) trained a classification algorithm to recognize stimulus-specific patterns in V1 using data from a perceptually-demanding orientation discrimination task, then used this classifier to successfully discriminate the identity of a




Figure 3-3: Event-related analysis of WM data. The mean response of the 60 most spatially selective voxels in V1 (collapsed across hemispheres) is shown as a function of memory condition and ROI. Data are collapsed across stimulus orientation (i.e., 45° vs. 135° trials) because no differences in mean response amplitude attributable to this factor were observed. Dashed vertical lines at 0 and 22s denote the onset of the sample and test stimuli, respectively. Shaded Regions represent ± 1 SEM.

remembered orientation during a separate WM task (see Harrison & Tong, 2009, for a similar demonstration). However, this analysis was restricted to cortical ROIs that corresponded to the location of the monitored or remembered stimulus. Here, we asked whether a qualitatively similar result would be observed in ROIs that were not retinotopically mapped to the position of the stimulus that was stored in WM (i.e., ROIs ipsilateral to the stimulus' location). To examine this possibility, we trained a classifier to recognize stimulus orientation using data from PM scans. Recall that the stimuli used in PM scans were identical to those used in the WM task. However, during PM scans, stimuli cycled on and off at 2Hz for the duration of each 15 second trial, and thus, WM

maintenance was not explicitly required. The classifier trained using the PM task was then used to infer the specific orientation that subjects were remembering during each trial of the WM task.

Figure 3-4 depicts the results of this analysis using the 75 most spatially-selective voxels from each cortical ROI. Data from one subject were excluded from this analysis due to 0% classification accuracy in each ROI across all tested pattern sizes. The findings reported here represent the remaining 16 subjects. However, a qualitatively similar pattern of results was also obtained when all subjects were included in the analysis. Collapsed across ROIs, this analysis revealed a successful generalization between patterns of activity observed during PM and WM maintenance, but only when active memory maintenance was required (two-tailed t-tests against chance, $p < .01$ and $p = .44$ for memory and no-memory conditions, respectively). Qualitatively similar results were observed across a range of pattern sizes (50 voxels). Thus, patterns of activity in a given ROI were similar during PM and WM maintenance. These findings lend further support to the hypothesis that WM involves the reiteration of a specific perceptual event. The same analysis failed to reveal consistent above chance classification in extrastriate regions V2v, V3v, hV4v, and V3a.

These findings reflect conditions in which the sample stimulus appeared in the same spatial location during the PM and WM tasks. We next asked whether the pattern of activity in a given ROI contralateral to the stimulus in the PM task could successfully predict the pattern of activity in the same ROI when it was ipsilateral to the sample in the WM task. Likewise, we also tested whether the pattern of activity in a given ROI ipsilateral to the stimulus in the PM task could successfully predict the pattern of activity

in the same ROI when it was contralateral to the sample in the WM task. These analyses failed to reveal consistent above-chance classification accuracy in any of the visual areas that we examined (two-tailed t-tests against chance, all $ps > .15$), suggesting that similarities between PM and WM maintenance may be restricted to situations where a monitored and remembered stimulus occupy the same spatial location.

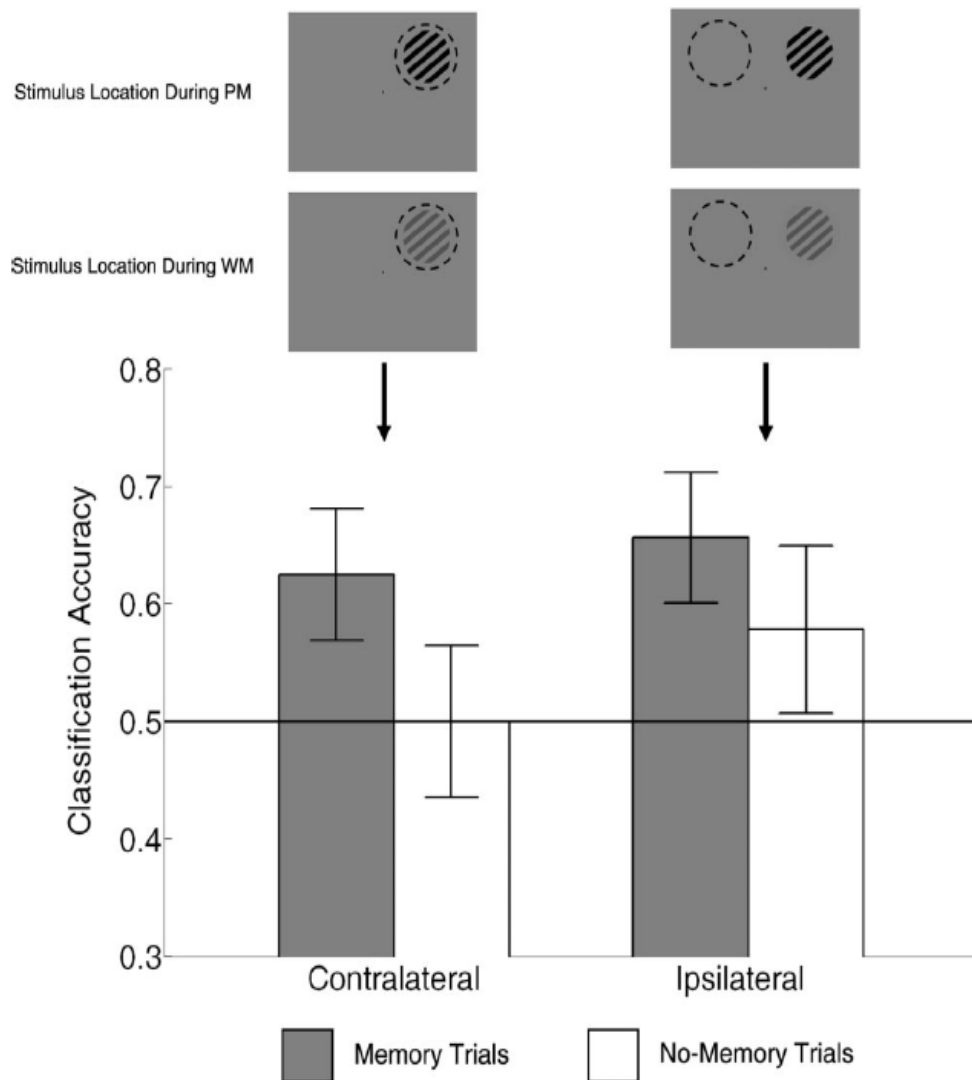


Figure 3-4: Global orientation-selective patterns are similar during PM and WM maintenance. A classification algorithm was trained to recognize orientation-selective patterns using the 75 most spatially selective voxels from each V1 ROI during PM (both left and right visual field trials were included in this analysis; top panels depict RVF trials for illustrative purposes). The algorithm was then used to infer the orientation of the same stimulus during each trial of the WM task. The horizontal line at 0.5 denotes chance classification accuracy. Above-chance classification accuracy was observed in ROIs both contralateral and ipsilateral to the sample stimulus, indicating that orientation-selective patterns are similar during PM and WM storage. Critically, this relationship was dependent on the observer's intent to store the stimulus' orientation during the WM task; we failed to observe above-chance classification accuracy on no-memory trials. Error bars depict ± 1 SEM.

DISCUSSION

An emerging perspective – informed by human neuroimaging and single-unit recordings in nonhuman primates – is that WM storage is mediated by sensory recruitment, or sustained activity in posterior cortical regions that encode memoranda (Miller et al., 1996; Awh & Jonides, 2001; Super et al., 2001; Jonides et al., 2005; Pasternak & Greenlee, 2005; Postle, 2006; D’Esposito, 2007). In this study, we attempted to determine the spatial extent of sensory recruitment during WM maintenance. Positional specificity effects in WM (e.g., Dill & Fahle, 1998; Zaksas, et al., 2001; Hollingworth 2006; 2007) motivate the hypothesis that sensory recruitment is spatially local. However, several studies have reported spatially global feature-based attention effects during perception (e.g., Treue & Martinez-Trujillo, 1999; Saenz, Buracas, & Boynton, 2002; Martinez-Trujillo & Treue, 2004; Bichot et al., 2005; Serences & Boynton, 2007a). Given known similarities between mechanisms of visual attention and working memory (e.g., Desimone, 1996; Rainer, Assad, & Miller, 1998; Awh & Jonides, 2001), these findings suggest that sensory recruitment may be spatially global. Here, subjects were required to remember the orientation of a grating presented in one visual hemifield. Using fMRI and MVPA, we found that patterns of activity in early visual areas (e.g., V1) both contralateral and ipsilateral to this stimulus discriminated its orientation (see Figure 3-2). These findings cannot be explained by subtle differences in head or eye movements. Moreover, under identical sensory conditions where WM maintenance was not explicitly required, these effects were abolished. Thus, above-chance classification observed on memory trials was not caused by a sustained sensory response to the sample

stimulus. Additionally, we found that patterns of activity in contralateral and ipsilateral ROIs were qualitatively similar during perceptual processing and WM maintenance, suggesting that similar mechanisms may support both processes (Figure 3-4). Taken together, these results suggest that sensory recruitment during WM maintenance is spatially global.

Several studies have reported successful classification of stimulus orientation in extrastriate visual areas during sensory encoding (Kamitani & Tong, 2005; Serences & Boynton, 2007b) and WM maintenance (Harrison & Tong, 2009). Likewise, studies of feature-based attention have documented spatially global effects in extrastriate visual areas (e.g., Treue & Martinez-Trujillo, 1999; Serences & Boynton, 2007a). In contrast, the spatially global WM effects reported here were confined to primary visual cortex. We suspect that this difference is due to the relatively low salience of our stimuli. For example, the large, high-contrast, and foveally-presented stimuli used by Harrison and Tong (2009) activated several hundred voxels within extrastriate visual areas V2-hV4v. In contrast, our small, peripheral stimuli produced a relatively weak response in only a few dozen voxels within each extrastriate visual area we examined (see Supplementary Table 1). It is possible that with a larger, sustained response, patterns of activity in each of these visual areas would support robust classification of the remembered orientation.

Our findings complement recent studies that have reported feature-specific patterns of activation in retinotopically organized visual areas during WM maintenance in the absence of sustained changes in overall response amplitude (e.g., Serences et al., 2009a; Harrison & Tong, 2009). In the current study, we observed a relative decrease in response amplitude during WM maintenance (relative to a condition where WM

maintenance was not required). One possibility is that WM maintenance involves the recruitment of only those neurons tuned to the remembered feature value, while neurons tuned to other values are suppressed (relatively speaking). This might lead to an overall decrease in response amplitude during WM storage. However, we emphasize that this account is purely speculative as changes in response amplitude have been inconsistent across different studies that have examined activity in primary visual cortex. For example, Serences et al. (2009a) failed to observe sustained increases in response amplitude during WM maintenance (see also Offen et al., 2009). Likewise, Harrison and Tong (2009) observed sustained increases in response amplitude during WM maintenance in a subset of their subjects, and no sustained changes in others. These studies, in concert with the current findings, raise questions about the extent to which sustained increases in response amplitude are diagnostic of an area's involvement in WM maintenance.

Our findings are consistent with a growing number of studies that have reported spatially global signals in retinotopic cortex during visual perception (e.g., Zaksas & Pasternak, 2005; Serences & Boynton, 2007a; Williams et al, 2008). In one example, Williams et al. (2008) required subjects to compare the category membership of two stimuli presented in the periphery of a visual display. Using fMRI and a multivariate analysis, these authors identified category-specific information in foveal retinotopic cortex, despite the fact that all stimuli were presented in spatial locations corresponding to cortical regions well outside of this area. Over several control experiments, the authors establish that these effects are likely due to feedback generated in anterior cortical regions. One difference between the task devised by Williams et al. and the WM task

used in this study is that in the former, all stimuli were present for the duration of each trial. However, we speculate that similar feedback mechanisms support the spatially global mnemonic representations observed in this study. One possibility is that this feedback directly facilitates stimulus-specific patterns of activity in both cortical hemispheres. Alternatively, feedback signals might be directed to a contralateral ROI and spread to ipsilateral ROIs via hard-wired cross-hemispheric connections between similarly tuned neurons in each hemisphere (for example, recent work has described a network of callosal axons linking cortical regions that represent the same orientation and spatial location of visual stimuli in cat visual areas 17 and 18; Rochefort et al., 2009; Schmidt et al., 2010). By this account, similarly tuned neurons in each cortical hemisphere are connected in a mutually excitatory manner; the efficacy of these connections might be modified by WM maintenance, giving rise to stimulus-specific modulations in the ipsilateral ROI. However, we believe this possibility unlikely given that patterns of activity in a given ROI during contralateral and ipsilateral trials are not necessarily identical (see Results). Finally, inhibitory connections between corresponding visual areas in each hemisphere may play a role in producing stimulus-specific activation patterns. This possibility is consistent with the lack of generalization across patterns of activation evoked by contralateral and ipsilateral memoranda in a given ROI.

An important question concerns the boundary conditions that govern the spatially global mnemonic representations observed in this experiment. Several studies have demonstrated that performance on a memory-limited change detection task is degraded when the sample and test stimuli can occupy different spatial locations (Foster & Khan, 1985; Dill & Fahle, 1998; Zaksas, et al., 2001). For example, Zaksas et al. (2001;

Experiment 1) trained monkeys to perform a change detection task while varying the spatial separation of the sample and test stimuli. These conditions were blocked, so it was always possible to infer the location of the test stimulus given the location of the sample. On some trials, the experimenters presented a random dot mask during the delay period; this stimulus could appear at the location of the sample stimulus, the impending test stimulus, or elsewhere in the display. Task performance was significantly impaired by the presentation of this mask, but only when it appeared at the location of the impending test stimulus. The selective interference produced by this stimulus suggests that memory for the random dot aperture was localized in the cortical regions that were retinotopically mapped to the position of the impending test stimulus. However, these findings do not necessarily preclude the existence of spatially global mnemonic representations similar to those observed in this study. As Zaksas et al note, the masking effects observed in this study may simply mean that the discrimination judgment required by the task was biased towards information stored in the cortical regions that processed the eventual test location, even if the relevant information was also stored in other cortical regions. It is therefore still possible that the relevant directional information was stored in a spatially global manner during this task.

Do the spatially global activation patterns observed in this study have functional consequences for WM maintenance? Posterior visual areas V1-hV4v are thought to represent information about feature properties such as orientation via population response profiles (e.g., Pouget et al., 2001; Pouget et al., 2003; Ma et al., 2006). One possibility is that the spatially global representations observed in this study act to enhance the precision of mnemonic representations (e.g., Awh, Barton, & Vogel, 2007; Zhang &

Luck, 2008; Barton, Ester, & Awh, 2009). For example, spatially global sensory recruitment would increase the number of neurons dedicated to representing a remembered feature attribute, which may in turn improve the efficiency or signal-to-noise ratio of population responses and thus enhance the precision of mnemonic representations. Alternatively, these global representations may be a passive result of diffuse feedback projections from higher cortical areas. Future research is needed to distinguish these possibilities.

In summary, our findings demonstrate that sensory recruitment during active WM maintenance is spatially global. We propose that spatially global gain modulations such as those observed in this study may serve to enhance the precision of mnemonic representations by recruiting additional sensory neurons that are not directly driven by the stimulus.

CHAPTER IV

GENERAL CONCLUSIONS

Converging evidence from multiple sources suggests that the storage of fine visual detail in WM is mediated by *sensory recruitment*, or sustained activity in cortical regions responsible for the sensory processing of memoranda. Here, I examined two untested predictions of this model. In Chapter II, I examined whether sustained patterns of activation observed in sensory cortices during WM storage contain information about specific featural attributes (e.g., color, orientation) of a remembered stimulus. Consistent with this possibility, functional magnetic resonance imaging (fMRI) and multivoxel pattern analysis (MVPA) revealed that sustained patterns of activation observed in sensory cortices (specifically, V1) discriminate the specific visual features that an observer is remembering. Moreover, these patterns of activation were qualitatively similar to those observed during the sensory processing of identical stimuli, suggesting that mnemonic representations in V1 are reasonable “copies” of those evoked during sensory processing. In Chapter III, I examined the hypothesis that humans opportunistically recruit all available sensory machinery in order to represent a specific feature value with maximal precision. Consistent with the results presented in Chapter II, patterns of activation observed in sensory cortical regions of interest (ROIs) retinotopically mapped (i.e., contralateral) to a remembered stimulus discriminated its orientation. In addition, patterns of activation observed in ipsilateral ROIs (i.e., those mapped to portions of empty space) also discriminated the orientation of a remembered stimulus. These results suggest that visual details are held in WM via a spatially global

recruitment of sensory cortex. This global recruitment may enhance memory precision by facilitating robust population coding of information.

In sum, the findings reported here are broadly consistent with a sensory recruitment model of WM. Before concluding, I briefly discuss a number of outstanding issues, including (1) the diagnostic power of sustained changes in BOLD amplitude as they relate to WM storage, (2) the conditions under which sensory recruitment might or might not be observed, and (3) the role of sensory recruitment in the short-term storage of information in non-visual sensory modalities.

Sustained Changes in the BOLD Signal as a Measure of WM Storage

Sustained changes in the amplitude of the blood-oxygen-level-dependent (BOLD) signal during the delay period of a WM task are typically interpreted as evidence that a given cortical region contributes to the short-term storage of information. However, although the experiments described in Chapters II and III revealed sustained stimulus-specific patterns of activation in primary visual cortex while subjects remembered the orientation or color of a stimulus, neither revealed any evidence for a sustained increase in the amplitude of the BOLD response during this interval. This finding raises a number of important questions regarding the diagnostic power of this measure (see Serences and Saproo, *in press*, for an in-depth discussion of this issue). Consider a case where a subject is holding a specific orientation value in WM. Single-unit recording studies in non-human primates suggest that directing attention to a specific feature value enhances the responses of cells that prefer the attended value and suppresses the responses of cells that prefer orthogonal values (relatively speaking; e.g., Martinez-Trujillo & Treue, 2004); this

pattern of activity is thought to represent the stored information (Paradiso, 1988; Pouget, Dayan, & Zemel, 2003; Sanger, 1996). According to the sensory recruitment hypothesis, the pattern of activity observed when an observer directs attention to a specific feature value should persist while an observer holds this information in WM. However, because each fMRI voxel integrates information from neurons that are more active (i.e., those tuned to the remembered orientation) with information from neurons that are less active, one might expect little overall change in the amplitude of the overall BOLD signal. Thus, the absence of a sustained change in the BOLD response during the delay period of a WM task should not be interpreted as definitive evidence that a given cortical region does not contribute to WM storage.

Under What Conditions Does Sensory Recruitment Occur?

Researchers who work in the domain of WM are (often painfully) aware that humans “opportunistically, automatically, recruit as many mental codes as are afforded by a stimulus when representing that stimulus in working memory” (Postle, 2006; p. 31). For example, a subject might choose to remember a color by assigning it a categorical label (e.g., “light red” or “dark green”) instead of attempting to store its exact hue. Alternately, a subject might choose to remember an array of spatial positions as the vertices of a single polygon. These “recoding” strategies can have a significant influence of memory performance, and researchers typically take great pains to discourage their use. Thus, the experiments described in Chapters II and III were designed with the assumption that sensory recruitment would be most evident when recoding strategies were discouraged. To this end, subjects were required to make very fine-grained

discriminations between sample and test stimuli (where the magnitude of sample-test changes were determined on a subject-by-subject basis using a titration procedure) in the hope that this would actively discourage the use of verbal labels or categorical coding.

Does sensory recruitment also occur when observers are allowed to use alternative mnemonic strategies, or in tasks that do not require the storage of detailed visual information (e.g., the change detection procedure described by Luck & Vogel, 1997)? To my knowledge, this issue has never been systematically investigated. However, from a purely theoretical perspective there is ample room for doubt on this point. As mentioned above, most contemporary neural models suggest that WM storage involves the opportunistic recruitment of different cortical regions to represent different types of information via different mnemonic strategies (e.g., Postle, 2006; D’Esposito, 2007). Thus, if an observer can effectively store visual information using a non-visual (e.g., verbal or categorical) code, then there is little reason to engage in the metabolically demanding task of recruiting specific portions of visual cortex in an effort to maintain a “veridical” stimulus representation.

Sensory Recruitment in Non-visual Sensory Modalities

The experiments described in Chapters II and III focused on sensory recruitment in the domain of visual WM. However, as noted in Chapter I, there is ample reason to suspect that sensory recruitment mediates the storage of fine detail for non-visual stimuli. For example, single-unit recordings in macaques have revealed sustained changes in the firing of neurons in the primary somatosensory cortex (S1) during the delay period of a haptic memory task (Zhou & Fuster, 1996). In a related study, Harris et al. (2002) asked

human subjects to compare two vibrotactile stimuli delivered to the same hand and separated by a brief delay. During the delay period, the authors delivered a single transcranial magnetic stimulation (TMS) pulse to regions of S1 contralateral or ipsilateral to the stimulated hand. Relative to pulses delivered to ipsilateral portions of S1, pulses delivered to portions of S1 contralateral to the stimulated hand had a deleterious effect on discrimination performance. This finding is consistent with the hypothesis that S1 contributes to the storage of vibrotactile information in WM². Other studies, however, have failed to find evidence for sustained changes in S1 responses during WM storage. For example, Romo and colleagues (e.g., Hernandez, Zainos, & Romo, 2000; Salinas, Hernandez, Sainos, & Romo, 2000) recorded from macaque S1 during the delay period of a task that required monkeys to discriminate between two vibrotactile stimuli. These authors observed transient changes in the responses of S1 neurons immediately following the presentation of the sample stimulus, but this activity quickly decayed during the subsequent delay interval. Harris et al. (2002) suggested that this result could reflect the fact that monkeys were extensively trained in the discrimination task prior to recording; this may allowed the monkeys to adopt behavioral strategies that do not rely on activity in S1. Further research is necessary to explore this possibility in detail.

Evidence for sensory recruitment has also been observed in auditory WM. In one example, Gottlieb, Vaadia, and Abeles (1989) recorded from the primary auditory cortex (A1) of macaques during the delay period of a pitch discrimination task. These authors

² Note that TMS pulses applied to regions of S1 ipsilateral to the stimulated hand had no effect on discrimination performance. This is inconsistent with a key finding of the experiment reported in Chapter III, where sensory recruitment was found in regions of visual cortex both contralateral and ipsilateral to a remembered stimulus. A likely explanation for this difference is anatomical: whereas regions of left and right primary visual cortex with similar spectral tuning are connected by cross-callosal fibers (e.g., Rochefort et al., 2009; Schmidt et al., 2010), hand representations in S1 are acallosal (Jones & Powell, 1969; Killackey et al., 1983).

observed sustained increases in the activity of A1 neurons selective for the first of the two tones that persisted throughout the delay period. Critically, this effect was drastically reduced (though not fully abolished) during an identical task that did not require the monkeys to compare the two tones. In a related study, Brechmann et al. (2007) asked human subjects to perform an n-back auditory WM task that required storage of frequency-modulated tones over short delays. Using fMRI, these authors observed sustained changes in bilateral regions of the planum temporal that predicted subjects' behavioral performance. Together, these results are consistent that A1 and other early auditory areas play an important role in the short-term storage of pitch information. However, it should be noted that more recent studies (e.g., Lemus, Hernandez, & Romo, 2009a; 2009b) have failed to identify sustained changes in the activity of A1 neurons during auditory delayed-match-to-sample tasks. In one example, monkeys were required to discriminate the pitches of two sequential acoustic "flutter" stimuli (rapid pulses of an acoustic stimulus) separated by a brief delay. Single-unit recordings failed to reveal any sustained changes in the firing rates of A1 neurons during the delay period of this task. At present, it is unclear what factor(s) might account for the discrepancies between these studies, though one possibility is that extensive training in the discrimination task prior to recording allowed monkeys to learn and adopt behavioral strategies that do not rely on the responses of A1 neurons.

Conclusions

An emerging perspective – informed by human neuroimaging and single-unit recordings in nonhuman primates – is that WM storage is mediated by sensory

recruitment, or sustained activity in posterior cortical regions that encode memoranda (Miller et al., 1996; Awh & Jonides, 2001; Super et al., 2001; Jonides et al., 2005; Pasternak & Greenlee, 2005; Postle, 2006; D’Esposito, 2007). In this dissertation, I examined two untested predictions of this perspective. In Chapter II, fMRI and MVPA were used to demonstrate that during a WM task that requires the storage of fine visual detail, sustained patterns of activation observed in sensory cortices (specifically, V1) discriminate the specific visual features that an observer is remembering. Moreover, these patterns of activation were qualitatively similar to those observed during the sensory processing of identical stimuli, suggesting that mnemonic representations in V1 are reasonable “copies” of those evoked during sensory processing. In Chapter III, similar methods were used to demonstrate that humans opportunistically recruit all available sensory machinery in order to represent a specific feature value with maximal precision. Specifically, patterns of activation observed in sensory cortical regions of interest (ROIs) retinotopically mapped (i.e., contralateral) to a remembered stimulus discriminated its orientation. In addition, patterns of activation observed in ipsilateral ROIs (i.e., those mapped to portions of empty space) discriminated the orientation of a stimulus that observers were storing in WM. Together, these findings lend additional support to the sensory recruitment view of WM.

REFERENCES CITED

- Alvarez GA, Cavanagh P (2004) The capacity of visual short-term memory is set both by visual information load and by number of objects. *Psychol Sci* 15:106-111.
- Anderson DE, Vogel EK, Awh E (2011) Precision in visual working memory reaches a stable plateau when individual item limits are exceeded. *J Neurosci* 31:1128-1138.
- Awh E, Barton B, Vogel EK (2007) Visual working memory represents a fixed number of items, regardless of complexity. *Psychol Sci* 18:622-628.
- Awh E, Jonides, J. (2001) Overlapping mechanisms of attention and spatial working memory. *Trends Cogn Sci* 5:119-126.
- Baddeley AD, Bressi S, Della Sala S, Logie R, Spinnler H (1991) The decline of working memory in Alzheimer's disease: A longitudinal study. *Brain* 114:2521-2542.
- Barton B, Ester EF, Awh E (2009) Discrete resource allocation in visual working memory. *J Exp Psychol Hum Percept Perform* 35:1359-1367.
- Bays PM, Husain M (2008) Dynamic shifts of limited working memory resources in human vision. *Science* 321:851-854.
- Bays PM, Catalao RFG, Husain M (2009) The precision of visual working memory is set by allocation of a shared resource. *J Vis* 9:1-11.
- Bays PM, Wu EY, Husain M (2011) Storage and binding of object features in visual working memory. *Neuropsychologia* 49:1622-1631.
- Bichot NP, Rossi AF, Desimone R (2005) Parallel and serial neural mechanisms for visual search in macaque area V4. *Science* 308:529-534.
- Bisley JW, Pasternak T (2000) The multiple roles of visual cortical areas MT/MST in remembering the direction of visual motion. *Cereb Cortex* 10:1053-1065.
- Bisley JW, Zaksas D, Pasternak T (2001) Microstimulation of cortical areas MT affects performance on a visual working memory task. *J Neurophysiol* 85:187-196.
- Bisley JW, Zaksas D, Droll JA, Pasternak T (2004) Activity of neurons in cortical area MT during a memory for motion task. *J Neurophysiol* 91:286-300.
- Boynton GM, Engel SA, Glover GH, Heeger DJ (1996) Linear systems analysis of functional magnetic resonance imaging in human V1. *J Neurosci* 16:4207-4221.
- Brainard DH (1997) The Psychophysics Toolbox. *Spat Vis* 10:433-436.

- Brechmann A, Gaschler-Markefski B, Sohr M, Yoneda K, Kaulisch T, Scheich H (2007) Working memory-specific activity in auditory cortex: Potential correlates of sequential processing and maintenance. *Cereb Cortex* 17:2544-2552.
- Channon S, Baker JE, Robertson MM (1993) Working memory in clinical depression: An experimental study. *Psychol Med* 23:87-91.
- Chelazzi L, Miller EK, Duncan J, Desimone R (1993) A neural basis for visual search in inferior temporal cortex. *Nature* 363:345-347.
- Conway ARA, Cowan N, Bunting MF, Theriault DJ, Minkoff RB (2002) A latent variable analysis of working memory capacity, short-term memory capacity, processing speed, and general fluid intelligence. *Intelligence* 30:163-183.
- Courtney SM, Ungerleider LG, Keil K, Haxby JV (1996) Object and spatial working memory activate separate neural systems in human cortex. *Cereb Cortex* 6:39-49.
- Courtney SM, Ungerleider LG, Keil K, Haxby JV (1997) Transient and sustained activity in a distributed neural system for human working memory. *Nature* 386:608-611.
- Daneman M, Carpenter PA (1980) Individual differences in working memory and reading. *J Mem Lang* 19:450-466
- D'Esposito M (2007) From cognitive to neural models of working memory. *Philos Trans R Soc Lond B Biol Sci* 362:761-772.
- Desimone R (1996) Neural mechanisms for visual memory and their role in attention. *Proc Natl Acad Sci USA* 93:13494-13499.
- Dill M, Fahle M (1998) Limited translation invariance of human visual pattern recognition. *Percept Psychophys* 60:65-81.
- Druzgal TJ, D'Esposito M (2001) Activity in fusiform face area modulated as a function of working memory load. *Cogn Brain Res* 10:355-364.
- Eng HY, Chen D, Jiang Y (2005) Visual working memory for simple and complex visual stimuli. *Psychon Bull Rev* 12:1127-1133.
- Engel SA, Rumelhart DE, Wandell BA, Lee AT, Glover GH., Chichilnisky EJ, Shadlen MN (1994) fMRI of human visual cortex. *Nature* 369:525.
- Engle RW, Tuholski SW, Laughlin JE, Conway A (1999) Working memory, short-term memory, and general fluid intelligence: A latent-variable approach. *J Exp Psychol Gen* 128:309-331.

- Ettlinger G (1990) "Object vision" and "spatial vision": The neuropsychological evidence for the distinction. *Cortex* 26:319-341.
- Foster DH, Khan JI (1985) Internal representations and operations in the visual comparison of transformed patterns: effects of pattern point-inversion, positional symmetry, and separation. *Biol Cybern* 51:305-312.
- Fukuda K, Awh E, Vogel EK (2010) Discrete capacity limits in visual working memory. *Curr Opin Neurobiol* 20:177-182.
- Fukuda K, Vogel EK, Mayr U, Awh E (2010) Quantity, not quality: The relationship between fluid intelligence and working memory capacity. *Psychon Bull Rev* 17:673-679.
- Fuster JM (1973) Unit activity in prefrontal cortex during delayed-response performance: Neuronal correlates of transient memory. *J Neurophysiol* 36:61-78.
- Fuster JM, Alexander GE (1971) Neuron activity related to short-term memory. *Science* 173:652-654.
- Gold JM, Carpenter C, Randolph C, Goldberg TE, Weinberger DR (1997) Auditory working memory and Wisconsin card sorting test performance in schizophrenia. *Arch Gen Psychiatry* 54:159-165.
- Gold JM, Wilk CM, McMahon RP, Buchanan R, Luck SJ (2003) Working memory for visual features and conjunctions in schizophrenia. *J Abnorm Psychol* 112:61-71.
- Goldman PS, Rosvold HE (1970) Localization of function within the dorsolateral prefrontal cortex of the rhesus monkey. *Exp Neurol* 27:291-304.
- Goldman-Rakic PS (1994) Working memory dysfunction in schizophrenia. *J Neuropsychiatry Clin Neurosci* 6:348-357.
- Gottlieb Y, Vaadia E, Abeles M (1989) Single unit activity in the auditory cortex of a monkey performing a short term memory task. *Exp Brain Res* 74:139-148.
- Gross CG (1963) A comparison of the effects of parietal and total lateral frontal lesions on test performance by monkeys. *J Com Physiol Psychol* 56:41-47.
- Harrison SA, Tong F (2009) Decoding reveals the contents of visual working memory in early visual areas. *Nature* 458:632-635.
- Haxby JV, Gobbini MI, Furey MI, Ishai A, Schouten JL, Pietrini P (2001) Distributed and overlapping representations of faces and objects in ventral temporal cortex. *Science* 293:2325-2430.

- Harris JA, Miniussi C, Harris IM, Diamond ME (2002) Transient storage of a tactile memory trace in primary somatosensory cortex. *J Neurosci* 22:8720-8725.
- Haynes JD, Rees G (2005) Predicting the orientation of invisible stimuli from activity in human primary visual cortex. *Nat Neurosci* 8:686-691.
- Hernandez A, Zainos A, Romo R (2000) Neuronal correlates of sensory discrimination in the somatosensory cortex. *Proc Natl Acad Sci USA* 97:6191-6196.
- Hollingworth A (2006) Scene and position specificity in visual memory for objects. *J Exp Psychol Hum Learn* 32:58-69.
- Hollingworth A (2007) Object-position binding in visual memory for natural scenes and object arrays. *J Exp Psychol Hum Percept Perform* 33:31-47.
- Irwin DE (1992) Memory for position and identity across eye movements. *J Exp Psychol Learn Mem Cogn* 18:307-317
- Johnson EN, Hawken MJ, Shapley R (2001) The spatial transformation of color in the primary visual cortex of the macaque monkey. *Nat Neurosci* 4:409-416.
- Jones EG, Powell TP (1969) Connections of the somatic sensory cortex of the rhesus monkey. II. Contralateral cortical connexions. *Brain* 92:717-730.
- Jonides J, Lacey SC, Nee DE (2005) Processes of working memory in mind and brain. *Curr Dir Psychol Sci* 14:2-5.
- Kamitani Y, Tong F (2005) Decoding the visual and subjective contents of the human brain. *Nat Neurosci* 8:1248-1254.
- Kamitani Y, Tong F (2006) Decoding seen and attended motion directions from activity in the human visual cortex. *Curr Biol* 16:1096-1102.
- Killackey HP, Hould HJ III, Cusick CG, Pons TP, Kaas JH (1983) The relation of corpus callosum connections to architectonic fields and body surface maps in sensorymotor cortex of new and old world monkeys. *J Comp Neurol* 219:384-419.
- Lemus L, Hernandez A, Romo R (2009a) Neural codes for perceptual discrimination of acoustic flutter in the primate auditory cortex. *Proc Natl Acad Sci USA* 106:9471-9476.
- Lemus L, Hernandez A, Romo R (2009b) Neural encoding of auditory discrimination in ventral premotor cortex. *Proc Natl Acad Sci USA* 106:14640-14645.

- Lepsien J, Nobre AC (2007) Attentional modulation of object representations in working memory. *Cereb Cortex* 17:2072-2083.
- Leventhal AG, Thompson KG, Liu D, Zhou Y, Ault SJ (1995) Concomitant sensitivity to orientation, direction, and color of cells in layers 2, 3, and 4 of monkey striate cortex. *J Neurosci* 15:1808-1818.
- Lisman JE, Idiart MA (1995) Storage of 7 ± 2 short-term memories in oscillatory subcycles. *Science* 267:1512-1515.
- Logothetis NK, Pauls J, Augath M, Trinath T, Oeltermann A (2001) Neurophysiological investigation of the basis of the fMRI signal. *Nature* 412:150-157.
- Luck SJ, Vogel EK (1997) The capacity of visual working memory for features and conjunctions. *Nature* 390:279-281.
- Ma WJ, Beck JM, Latham PE, Pouget A (2006) Bayesian inference with probabilistic population codes. *Nat Neurosci* 9:1432-1438.
- Mangussen S, Greenlee MW, Asplund R, Dyrnes S (1991) Stimulus-specific mechanisms of visual short-term memory. *Vision Res* 31:1213-1219.
- Martinez-Trujillo JC, Treue S (2004) Feature-based attention increase the selectivity of population responses in primate visual cortex. *Curr Biol* 13:655-662.
- Miller EK, Li L, Desimone R (1993) Activity of neurons in anterior inferior temporal cortex during a short-term memory task. *J Neurosci* 13:1460-1478.
- Miller EK, Erickson CA, Desimone R (1996) Neural mechanisms of visual working memory in prefrontal cortex of the macaque. *J Neurosci* 16:5154-5167.
- Niki H, Watanabe M (1976) Prefrontal unit activity and delayed response: relation to cue location versus direction of response. *Brain Res* 105: 79-88.
- Norman KA, Polyn SM, Detre GJ, Haxby JV (2006) Beyond mind-reading: multi-voxel pattern analysis of fMRI data. *Trends Cogn Sci* 10: 424-430.
- Offen S, Schluppeck D, Heeger DJ (2009) The role of early visual cortex in visual short-term memory and visual attention. *Vision Res* 49:1352-1362.
- Olivers CNL, Meijer F, Theeuwes J (2006) Feature-based memory-driven attentional capture: Visual working memory content affects visual attention. *J Exp Psychol Hum Percept Perform* 32:1243-1265.

- Owen, AM, Iddon JL, Hodges JR, Summers BA, Robbins TW (1997) Spatial and non-spatial working memory at different stages of Parkinson's disease. *Neuropsychologia* 35:519-532.
- Paradiso MA (1998) A theory for the use of visual orientation information which exploits the columnar structure of striate cortex. *Biol Cybern* 58:35-49.
- Pasternak T, Greenlee M (2005) Working memory in primate sensory systems. *Nat Rev Neurosci* 6:97-107.
- Peelen MV, Downing PE (2007) Using multi-voxel pattern analysis of fMRI data to interpret overlapping functional activations. *Trends Cogn Sci* 11:4-5.
- Pelli DG (1997) The VideoToolbox software for visual psychophysics: Transforming numbers into movies. *Spat Vis* 10:437-442.
- Postle BR (2006) Working memory as an emergent property of the mind and brain. *Neuroscience* 139: 23-38.
- Pouget A, Dayan P, Zemel R (2000) Information processing with population codes. *Nat Rev Neurosci* 1:125-132.
- Pouget A, Dayan P, Zemel R (2003) Inference and computation with population codes. *Annu Rev Neurosci* 26:381-410.
- Raffone A, Wolters G (2001) A cortical mechanism for binding in visual working memory. *J Cogn Neurosci* 13:766-785.
- Rainer G, Asaad WF, Miller EK (1998) Selective representation of relevant information by neurons in the primate prefrontal cortex. *Nature* 393:577-579.
- Ranganath C, Cohen MX, Dam C, D'Esposito M (2004) Inferior temporal, prefrontal, and hippocampal contributions to visual working memory maintenance and associative memory retrieval. *J Neurosci* 24:3917-3925.
- Rocheffort N, Buzas P, Quenech'du N, Koza A, Eysel UT, Milleret C, Kisvarday ZF (2009) Functional selectivity of interhemispheric connections in cat visual cortex. *Cereb Cortex* 19:2451-2465.
- Romo R, Hernandez A, Zainos A, Lemus L, Brody CD (2002). Neuronal correlates of decision making in secondary somatosensory cortex. *Nat Neurosci* 5:1217-1225.
- Rouder JN, Morey RD, Cowan N, Zwilling CE, Morey CC, Pratte MS (2008) An assessment of fixed-capacity models of visual working memory. *Proc Natl Acad Sci USA* 105:5975-5979.

- Saenz M, Buracas GT, Boynton GM (2002) Global effects of feature-based attention in human visual cortex. *Nature Neurosci* 5: 631-632.
- Salinas E, Hernandez A, Zainos A, Romo R (2000) Periodicity and firing rate as candidate neural codes for the frequency of vibrotactile stimuli. *J Neurosci* 20:5503-5515.
- Sanger TD (1996) Probability density estimation for the interpretation of neural population codes. *J Neurophysiol* 76:2790-2793.
- Schmidt, KE, Lomber SG, Innocenti GM (2010) Specificity of neuronal responses in primary visual cortex is modulated by interhemispheric corticocortical input. *Cereb Cortex* 20:2776-2786.
- Serences JT, Boynton GM (2007a) Feature-based attentional modulations in the absence of direct visual stimulation. *Neuron* 55:301-312.
- Serences JT, Boynton, GM (2007b) The representation of behavioral choice for motion in human visual cortex. *J Neurosci* 27:12893-12899.
- Serences JT, Ester EF, Vogel EK, Awh E (2009a) Stimulus-specific delay activity in human primary visual cortex. *Psychol Sci* 20:207-214.
- Serences JT, Saproo S, Scolari M, Ho T, Muftuler LT (2009b) Estimating the influence of attention on population codes in human visual cortex using voxel-based tuning functions. *Neuroimage* 44:223-231.
- Serences JT, Saproo S (in press) Computational advances towards linking BOLD and behavior. *Neuropsychologia*.
- Sereno MI, Dale AM, Reppas JB, Kwong KK, Belliveau JW, Brady TJ, et al. (1995) Borders of multiple visual areas in humans revealed by functional magnetic resonance imaging. *Science* 268:889-893.
- Siegel M, Warden MR, Miller EK (2009) Phase-dependent neuronal coding of objects in short-term memory. *Proc Natl Acad Sci USA* 106:21341-21346.
- Sincich LC, Horton JC (2005) The circuitry of V1 and V2: Integration of color, form, and motion. *Annu Rev Neurosci* 28:303-326.
- Solomon SG, Lennie P (2007) The machinery of colour vision. *Nat Rev Neurosci* 8:276-286.
- Sperling G (1960) The information available in brief visual presentations. *Psychol Monogr* 74:1-29.

- Stokes M, Thompson R, Cusack R, Duncan J (2009) Top-down activation of shape-specific population codes in visual cortex during mental imagery. *J Neurosci* 29:1565-1572.
- Super H, Spekreijse H, Lamme VAF (2001) A neural correlate of working memory in the monkey primary visual cortex. *Science* 293:120-124.
- Todd JJ, Marois R (2004) Capacity limit of visual short-term memory in human posterior parietal cortex. *Nature* 428:751-754.
- Treue S, Martinez -Trujillo JC (1999) Feature-based attention influences motion processing gain in macaque visual cortex. *Nature* 382:539-541.
- Ungerleider LG, Mishkin M (1982) Two cortical visual systems. In *Analysis of Visual Behavior* (J Ingle, MA Goodale, RJW Mansfield, eds), pp. 549-586. Cambridge, MA: MIT Press.
- Vogel EK, Machizawa MG (2004) Neural activity predicts individual differences in visual working memory capacity. *Nature* 428:748-751.
- Vogel EK, McCollough AW, Machizawa MG (2005) Neural measures reveal individual differences in controlling access to working memory. *Nature* 438:500-503.
- Warren JM, Akert K, eds. (1964) *The frontal granular cortex and behavior*. New York: McGraw-Hill Book Company.
- Williams MA, Baker CI, Op de Beeck HP, Shim WM, Dang S, Triantafyllou C, Kanwisher N (2008) Feedback of visual object information to foveal retinotopic cortex. *Nat Neurosci* 11:1439-1445.
- Wilken P, Ma WJ (2004) A detection theory account of change detection. *J Vis* 4(12).
- Woodman GF, Vogel EK (2008) Selective storage and maintenance of an object's features in working memory. *Psychon Bull Rev* 15:223-229.
- Xu Y, Chun MM (2006) Dissociable neural mechanisms supporting visual short-term memory for objects. *Nature* 440:233-235.
- Zaksas D, Bisley JW, Pasternak T (2001) Motion information is spatially localized in a visual working-memory task. *J Neurophysiol* 86:912-921.
- Zaksas D, Pasternak T (2006) Direction signals in the prefrontal cortex and in area MT during a working memory for visual motion task. *J Neurosci* 26:11726-11742.
- Zhang WW, Luck SJ (2008) Discrete fixed-resolution representations in visual working memory. *Nature* 453:233-235.

Zhou YD, Fuster JM (1996) Mnemonic neuronal activity in somatosensory cortex. Proc Natl Acad Sci USA 93:10533-10537.