Temporally Dissociable Mechanisms of Spatial, Feature, and Motor Selection during Working Memory–guided Behavior

Edward Ester and Rachel Weese

Abstract

■ Working memory (WM) is a capacity- and duration-limited system that forms a temporal bridge between fleeting sensory phenomena and possible actions. But how are the contents of WM used to guide behavior? A recent high-profile study reported evidence for simultaneous access to WM content and linked motor plans during WM-guided behavior, challenging serial models where task-relevant WM content is first selected and then mapped on to a task-relevant motor response. However, the task used in that study was not optimized to distinguish the selection of spatial versus nonspatial visual information stored in memory, nor to distinguish whether or how the chronometry of selecting nonspatial visual information stored in memory might differ from the selection of linked motor plans. Here, we revisited the chronometry of spatial, feature, and motor selection during WM-guided behavior using a task optimized to disentangle these processes. Concurrent EEG and eye position recordings revealed clear evidence for temporally dissociable spatial, feature, and motor selection during this task. Thus, our data reveal the existence of multiple WM selection mechanisms that belie conceptualizations of WM-guided behavior based on purely serial or parallel visuomotor processing. ■

INTRODUCTION

A fundamental purpose of any memory system is to inform future decisions and actions. This is particularly important in the case of working memory (WM), a duration- and capacity-limited system that forms a temporal bridge between fleeting sensory phenomena and possible actions. Recent theoretical conceptualizations of WM have begun to emphasize the action-oriented nature of this system (e.g., van Ede & Nobre, 2023; Heuer, Ohl, & Rolfs, 2020; Olivers & Roelfsema, 2020), and recent empirical findings suggest that behavioral (González-García, Formica, Liefooghe, & Brass, 2020; Ohl & Rolfs, 2020; Heuer & Schubo, 2017), circuit-level (Pho, Goard, Woodson, Crawford, & Sur, 2018), and systems-level (Boettcher, Gresch, Nobre, & van Ede, 2021; Galero-Salas et al., 2021; Rac-Lubashevsky & Frank, 2021; van Ede, Chekroud, Stokes, & Nobre, 2019; Chatham, Frank, & Badre, 2014) mechanisms of WM storage and action planning are tightly interwoven.

Recent studies suggest that human observers can store multiple stimulus–response mappings in WM (see van Ede & Nobre, 2023, for a review) and that behaviorally relevant WM content can be selected in parallel with required actions. In one high-profile example, van Ede, Chekroud, Stokes, et al. (2019) required participants to remember the orientations of two colored bars over a short delay. A color cue presented at the end of the delay informed participants which item would be probed for report, whereas the vertical tilt of the color-cue-matching bar informed participants which hand should be used for orientation recall (with clockwise and anticlockwise tilted bars requiring right- and left-hand responses, respectively). Cleverly, van Ede and colleagues independently manipulated the physical location of the probed bar (i.e., left vs. right visual hemifield) and the tilt of the probed bar (i.e., requiring a left- vs. right-hand response), which allowed them to distinguish the selection of visual and motor information via lateralized signals measured in contemporaneous EEG recordings. Surprisingly, EEG signals associated with the selection of visual and motor information had nearly identical time courses, suggesting that participants were able to select the task-relevant WM content and the appropriate response in parallel.

On the one hand, the findings reported by van Ede, Chekroud, Stokes, et al. (2019) seem to challenge classic conceptualizations of human memory performance based on sequential processing stages where relevant information is first selected and then mapped onto appropriate outputs (e.g., Meyer, Osman, Irwin, & Yantis, 1988; Donders, 1969; Sternberg, 1969). On the other hand, two aspects of the study performed by van Ede, Chekroud, Stokes, et al. (2019) undermine this challenge. First, the electrophysiological signal that these authors used to track the selection of visual information stored in WM—lateralized alpha-band activity—is known to index covert spatial attention (e.g., Klimesch, 2012). Although van Ede and colleagues presented to-be-remembered stimuli in opposite visual hemifields (i.e., one bar appeared in the left visual field and another in the right visual field), participants

University of Nevada, Reno

Journal of Cognitive Neuroscience 35:12, pp. 2014–2027
https://doi.org/10.1162/jocn_a_02061

were always probed to report a specific bar via a color cue and by adjusting a probe stimulus presented at fixation. Thus, participants could solve the task by storing color orientation bindings independent of where these bindings appeared. Although there is growing evidence that human observers automatically store spatial information in WM and use spatial information to select behaviorally relevant WM content (e.g., Groen, Dekker, Knapen, & Silson, 2022; Foster, Bsales, Jaffe, & Awh, 2017), this leaves open the question of how nonspatial features are selected alongside motor plans during WM-guided behavior.

Second, and more importantly, the design used by van Ede, Chekroud, Stokes, and Nobre (2019) confounded the selection of motor information with the selection of task-relevant feature information. Specifically, the vertical tilt of the retrospectively cued bar (i.e., clockwise vs. anticlockwise from vertical) instructed participants which hand should be used for recall (right or left, respectively). Thus, changes in lateralized beta-band EEG activity that these authors ascribed to motor selection could instead reflect feature selection (i.e., the clockwise or anticlockwise tilt of the to-be-recalled bar). Likewise, van Ede, Chekroud, Stokes, et al. (2019) yoked response hand with probe rotation direction during recall. That is, right-hand responses always produced clockwise rotations of the probe bar, whereas left-hand responses always produced anticlockwise rotations of the probe bar. Thus, the lateralized beta-band signals that these authors ascribed to motor planning could instead reflect a recall strategy based on feature information (e.g., "the probe stimulus must be rotated clockwise by 30°") rather than motor planning per se. These factors, coupled with studies linking EEG beta synchrony with content-specific WM storage (summarized in Spitzer & Haegens, 2017) and studies linking beta local field potentials in nonhuman primates with top–down control over the contents of WM (summarized in Miller, Lundqvist, & Bastos, 2018) motivate a fresh look at the role of beta oscillations in the selection of remembered feature and/or motor information.

In this study, we sought further clarity on the chronometry of spatial, feature, and motor selection during WM-guided behavior. We recorded EEG and eye position data while human volunteers performed a retrospectively cued orientation recall task (see Figure 1 for a schematic). Following earlier work (van Ede, Chekroud, Stokes, et al., 2019), we manipulated the physical location of the retrospectively cued orientation (i.e., left vs. right visual field) independently of response hand and used lateralized alpha and beta power to track the selection of spatial and motor information, respectively. Unlike prior work, we decoupled the selection of spatial, feature, and motor selection by (a) using a color cue to inform participants which item stored in WM was task-relevant and which hand should be used for recall, (b) delaying orientation recall until 1 sec after cue onset, and (c) using multivariate decoding analyses to track the selection of task-relevant orientation information stored in WM. To preview our findings, spectral analyses of EEG data and concurrently recorded eye position data time-locked to retrocue onset revealed evidence for the simultaneous selection of spatial and motor information following the onset of an informative retrocue, replicating van Ede, Chekroud, Stokes, et al. (2019). Critically, above-chance decoding of stimulus orientation was delayed until the onset of the response display, demonstrating that feature selection is temporally decoupled from spatial and motor selection. These findings challenge models of WM-guided behavior based on purely serial or parallel selection of visual and motor information and instead point to the existence of multiple temporally distinct selection mechanisms.

METHODS

Sample Size Justification

Before data collection, an a priori sample size of 32 human adult volunteers (both sexes) was selected. This choice was based on published empirical findings and effect sizes (e.g., van Ede, Chekroud, Stokes, et al., 2019).

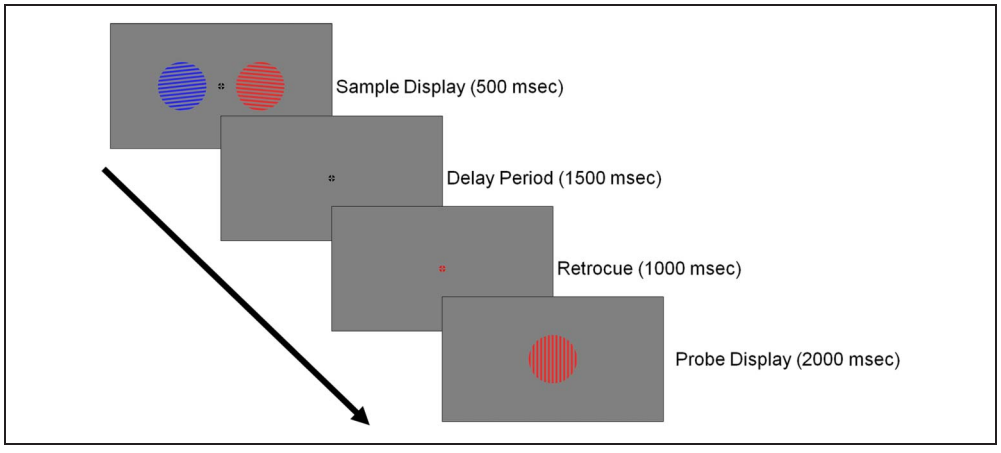

Figure 1. Experimental task. Participants remembered the orientations of two gratings over a blank delay and then recalled the orientation of a retrospectively cued grating. See text for details.

Research Participants

Thirty-two volunteers from the University of Nevada, Reno, participated in this experiment. Each participant completed a single 2.5-hr testing session in exchange for monetary remuneration ($15/hr). All volunteers self-reported normal or corrected-to-normal visual acuity and gave both written and oral informed consent before enrolling in the study. All study procedures were approved by the local institutional review board. Data from one participant who completed the study were excluded from final analyses because of chance-level task performance (i.e., the participant used the correct response hand on only 52% of trials); thus, the findings reported here reflect the remaining 31 participants.

Testing Environment

Participants were seated in a dimly lit room for the duration of testing. Visual displays were generated in MATLAB and rendered on a 27-in. LCD monitor (1920 × 1080 resolution) cycling at 240 Hz using Psychtoolbox-3 (Kleiner, Brainard, & Pelli, 2007). Participants were seated approximately 80 cm from the display (head position was not constrained). Participants responded via buttons on a standard U.S. computer keyboard.

Visuomotor Recall Task

A task schematic is shown in Figure 1. Each trial began with a 500-msec sample display containing two colored gratings (radius 4.5 degrees visual angle [DVA] from a viewing distance of 80 cm; 1.0 cycles/DVA) presented 7.0 DVA to the left and right of a central fixation point. The orientation of each grating on each trial was randomly and independently sampled (with replacement) from the set [20°, 40°, 60°, 80°, 120°, 140°, 160°), and a small amount of angular jitter (±1–5°) was added to each orientation on each trial to discourage verbal coding. Stimulus orientations were counterbalanced across the entire experiment (though not necessarily within a single block of trials). The sample display was followed by a 1500-msec blank display and a 1000-msec retrocue display where the fixation cross changed colors from black to either blue or red. The color of the retrospective cue instructed participants which grating would be probed at the end of the trial (i.e., blue or red) and which hand should be used to make a behavioral response (left or right). Cue color–response mappings (i.e., red cue/left hand vs. red cue/right hand) were counterbalanced across participants.

Each trial ended with a 2000-msec probe display containing a vertically oriented grating. The color of the probe grating always matched the color of the retrospectively cued grating. Participants were required to recall the orientation of the retrospectively cued grating by adjusting the probe to match the corresponding sample via the cued response hand. If the trial required a left-hand response, then participants were instructed to use the "Z" and "X" keys on a computer keyboard to rotate the probe grating counterclockwise versus clockwise, respectively. Conversely, if the trial required a right-hand response, then participants used the ">" and "?" keys to rotate the probe grating counterclockwise versus clockwise. We measured whether participants responded using the correct hand relative to their first keypress. For example, if a trial required a left-hand response and participants initially pressed either the ">" or "?" keys, the trial was counted as incorrect. Participants were free to adjust the probe stimulus for the full 2000-msec response period and instructed that if they were satisfied with their response before the end of the probe period, they should refrain from pressing any additional buttons. Participants completed three (n = 1), seven (n = 2), or eight (n = 28) blocks of 72 trials in this task.

Operationalization of Spatial, Motor, and Feature Selection

Following earlier work (van Ede, Chekroud, Stokes, et al., 2019), we tracked spatial and motor selection through analyses of lateralized occipitoparietal alpha power (8–13 Hz) and lateralized frontocentral beta power (15–30 Hz), respectively. Our strategy was based on observations that covert shifts of attention into one visual hemifield (e.g., left) produce a reduction in alpha power over contralateral occipitoparietal electrode sites (e.g., Keefe & Stöermer, 2021; Klimesch, 2012) that preparing a response using one hand (e.g., left) produces a reduction in beta power over contralateral frontocentral electrode sites (Kaiser, Birbaumer, & Lutzenberger, 2001). Thus, the timing(s) of changes in lateralized alpha and beta power can be used to track covert shifts of attention toward the visual hemifield containing the retrospectively cued grating (i.e., spatial selection) and the preparation of a motor response using the task-appropriate hand (i.e., motor selection), respectively. Finally, we operationalized feature selection through analyses of multivariate orientation decoding performance. This strategy was based on observations that directing attention to a cue-matching item stored in WM results in an increase in multivariate decoding performance for that item (e.g., Ester & Pytel, 2023; Ester, Nouri, & Rodriguez, 2018; LaRocque, Lewis-Peacock, Drysdale, Oberauer, & Postle, 2013). Thus, the timing(s) of changes in multivariate decoding performance can be used to track the selection of retrospectively cued feature information.

EEG Acquisition and Preprocessing

Continuous EEG was recorded from 63 scalp electrodes using a BrainProducts actiCHamp system. Online recordings were referenced to the left mastoid (10–20 site TP9) and digitized at 1 kHz. The following offline preprocessing steps were applied, in order: (1) resampling from 1 kHz to

500 Hz, (2) high-pass filtering (0.5 Hz using zero-phase forward- and reverse-finite impulse response filters as implemented by EEGLAB software extensions; Delorme & Makeig, 2004), (3) identification and reconstruction of noisy electrodes and epochs via artifact subspace reconstruction (implemented via EEGLAB; Chang, Hsu, Pion-Tonachini, & Jung, 2020), (4) re-referencing to the average response of all electrodes, (5) epoching from −1.0 to +5.0 sec relative to the start of each trial, (6) detection and removal of oculomotor and motor artifacts via independent components analysis and automated EEGLAB artifact detection functions, and (7) application of a surface Laplacian to remove low spatial frequency components from the signal (Perrin, Pernier, Bertrand, & Echallier, 1989). Omitting or altering any of these steps (e.g., re-referencing to the algebraic mean of the left and right mastoids or omitting artifact subspace reconstruction) had no qualitative impact on any of the findings reported in the article.

Eye-tracking Data Acquisition and Preprocessing

We obtained high-quality binocular eye position data for 18 of 32 participants enrolled in the study. Eye position data were acquired using an SR Research Eyelink 1000 Plus infrared eye tracker operating in remote (i.e., "head-free") mode and digitized at 500 Hz. The eye tracker was calibrated 2–3 times per testing session using a standard 9-point grid included with the tracker hardware. Eye position data were filtered for blinks and horizontal eye movements >2 DVA and then epoched from −1000 to +2000 msec relative to retrocue onset. No other preprocessing steps were applied.

Spectral EEG Analyses

Spectral analyses focused on occipitoparietal and frontocentral 10–20 electrode site pairs O1/2, PO3/4, PO7/8, C1/C2, and C3/C4; analyses were restricted to a period spanning −500 to +2500 around retrocue onset. We extracted broadband spectral power from each electrode on each trial using a short-time Fourier transform (STFT) with a frequency range of 1–40 Hz (in 0.25-Hz steps) using a 200-msec sliding window and 2-msec step size. Power was computed by squaring the absolute value of the complex Fourier coefficients within each STFT window. To quantify changes in lateralized activity during visual selection, we sorted power estimates at each occipitoparietal electrode site by the location of the retrospectively cued stimulus, that is, contralateral versus ipsilateral hemifield, and expressed this difference as a normalized percentage:

$$L_i = 100 \times \frac{(c_i - s_i)}{(c_i + s_i)}$$

where c_i and s_i are the trial-averaged responses of electrode i when the retrospectively cued grating appeared in the contralateral or ipsilateral visual field, respectively. Lateralization estimates were pooled across occipitoparietal electrode sites, yielding a single time frequency matrix of power estimates per participant. An identical approach was used to quantify changes in lateralized activity during motor selection, with the exceptions that (a) analyses focused exclusively frontocentral electrode sites and (b) trial-wise power estimates were sorted by the retrospectively cued response hand rather than stimulus location.

To generate topographical maps of alpha- and beta-band lateralization, we repeated this analysis for every scalp electrode and averaged power estimates over 8–13 Hz ("alpha-band" activity) and 15–30 Hz ("beta-band" activity) over a period of 500–1000 msec. These frequency bands were chosen a priori based on commonly reported values in the literature (e.g., Klimesch, 2012), and this specific temporal window was chosen a priori based on prior reports, suggesting that it takes human observers 300–500 msec to process and respond to a retrospective memory cue (e.g., Souza, Rerko, & Oberauer, 2014).

To visualize changes in occipitoparietal alpha power, frontocentral mu-alpha power, and frontocentral beta power, we extracted and averaged lateralization estimates from 8 to 13 Hz over electrode sites O1/2, PO3/4, and PO7/8, lateralization estimates from 8 to 13 Hz over electrode sites C1/2, C3/4, and lateralization estimates from 15 to 30 Hz over electrode sites C1/2, C3/4. Statistical analyses of lateralization were performed using nonparametric sign permutation tests with temporal cluster-based correction (see Statistical Comparisons section).

Quantifying Eye Position Biases

Preprocessed eye position data were sorted by the location of the retrospectively cued stimulus (i.e., left vs. right visual field). Because stimuli were rendered along the vertical meridian, we restricted our analyses to horizontal eye position measurements. Following earlier work (van Ede, Checkroud, & Nobre, 2019), we computed a normalized measure of gaze bias by converting pixelwise recordings into a percentage along an axis extending from fixation to the center of each visual stimulus. Thus, a ±100% gaze bias would result from participants foveating the center of one stimulus, whereas a 0% gaze bias would result from participants perfectly holding fixation. Gaze biases during right visual field trials were sign reversed and averaged with biases from left visual field trials, yielding a single gaze position bias time course per participant.

Orientation Decoding

Orientation decoding analyses were performed on broadband (0.5 Hz+) signals measured at occipitopartietal electrode sites O1/2, PO3/4, and PO7/8 and frontocentral electrode sites C1/2 and C3/4. Decoding performance was computed separately for each time sample using 10-fold

cross-validation. During each cross-validation fold, we designated 90% of available trials as a training data set and the remaining 10% of trials as a test data set, taking care to ensure that the training data set contained an equal number of observations for each orientation. Decoding performance was estimated using the multivariate distance between EEG activity patterns associated with memory for specific experimental conditions (Wolff, Jochim, Akyurek, & Stokes, 2017). We computed the Mahalanobis distance between trial-wise activation patterns in each test data set with position-specific activation patterns in the corresponding test data set, yielding a set of distance estimates. If scalp activation patterns contain information about stimulus orientation, then distance estimates should be smallest when comparing patterns associated with memory for similar or identical orientations in the training and test data sets and largest when comparing orthogonal orientations. Trial-wise distance functions were averaged and sign-reversed for interpretability. Decoding performance was estimated by convolving time point-wise distance functions with a cosine function, yielding a metric where chance decoding performance is equal to 0. Decoding results from each training and test data set pair were averaged (thus ensuring the internal reliability of our approach), yielding a single decoding estimate per participant and time point. To facilitate interpretability, orientation decoding time series were smoothed with a 100-msec Gaussian kernel. We verified that smoothing did not qualitatively affect any of the findings reported in the article.

Statistical Comparisons

Statistical comparisons were based on nonparametric signed randomization tests (Maris & Oostenveld, 2007). Unless otherwise specified, each test we performed assumes a null statistic of 0 (i.e., no difference in alpha- or beta-band lateralization, or no difference in observed vs. chance decoding performance). We therefore generated null distributions by randomly relabeling each participant's data with 50% probability and averaging the data across participants. This step was repeated 10,000 times, yielding a 10,000-element null distribution for each time point. Finally, we implemented a cluster-based permutation test with cluster-forming and cluster size thresholds of $p < .05$ (two-tailed) to evaluate observed differences with respect to the null distribution while accounting for signal autocorrelation.

RESULTS

We recorded EEG while 31 human volunteers performed a retrospectively cued orientation memory task (Figure 1). Participants remembered the orientations of two gratings over a blank delay. At 1.5 sec after encoding, a 100% valid color cue indicated which of the two gratings would be probed for report at the end of the trial and which hand should be used for recall (i.e., left or right; response hand/cue color mappings were counterbalanced across participants). A response display containing a vertically oriented grating was presented 1.0 sec later; participants adjusted the orientation of the probe to match the orientation of the cued grating using the cued response hand. Participants performed this task well, recalling the orientation of the probed grating with an average (± 1 *SEM*) absolute error of 13.6° ± 0.75° and responding with the correct hand on 93.8% ± 0.89% of trials. The average RT (i.e., elapsed time between the onset of the probe display and the participant's initial button press) was 451 ± 14.9 msec.

EEG Signatures of Spatial and Motor Selection

Following earlier work (van Ede, Chekroud, Stokes, et al. 2019), we first identified EEG signals that co-varied with the physical location of the retrospectively cued stimulus (i.e., left vs. right visual hemifield) and EEG signals that co-varied with the retrospectively cued response hand (i.e., left vs. right). Because stimulus location and response hand were independently manipulated over trials (i.e., a stimulus requiring a right-hand response was equally likely to appear in the left vs. right visual field), we were able to characterize the selection of both attributes in trial-averaged EEG data independently of one another and any nuisance effects (e.g., volume conduction or other signal mixing). We therefore attributed patterns of neural activity that covaried with stimulus location to the selection of spatial information stored in WM and patterns of neural activity that covaried with response hand to the selection of motor information stored in WM.

Spatial selection was associated with a robust but transient decrease in 8–13 Hz alpha power over electrode sites contralateral to the location of the cued stimulus (Figure 2A). This modulation was strongest over occipito-parietal electrode sites (Figure 2A, inset), consistent with prior reports linking changes in lateralized alpha power to shifts of attention in perception and WM (Klimesch, 2012). Conversely, motor selection was associated with a robust and sustained decrease in 8–13 Hz mu-alpha power and 15–30 Hz mu-beta power over electrode sites contralateral to the cued response hand (Figure 2B). This modulation was strongest over frontocentral electrode sites (Figure 2B, inset), consistent with prior reports linking changes in lateralized mu-alpha and beta power to response preparation and execution. Direct comparisons of EEG signals associated with visual and motor selection revealed transient visual selection lasting approximately 250–850 msec and sustained motor selection lasting 250–3000+ msec after retrocue onset (Figure 2C). However, neither the onset, magnitude, nor duration of spatial or motor selection during the cue-to-probe interval varied as a function of participants' RTs (i.e., time to first keypress

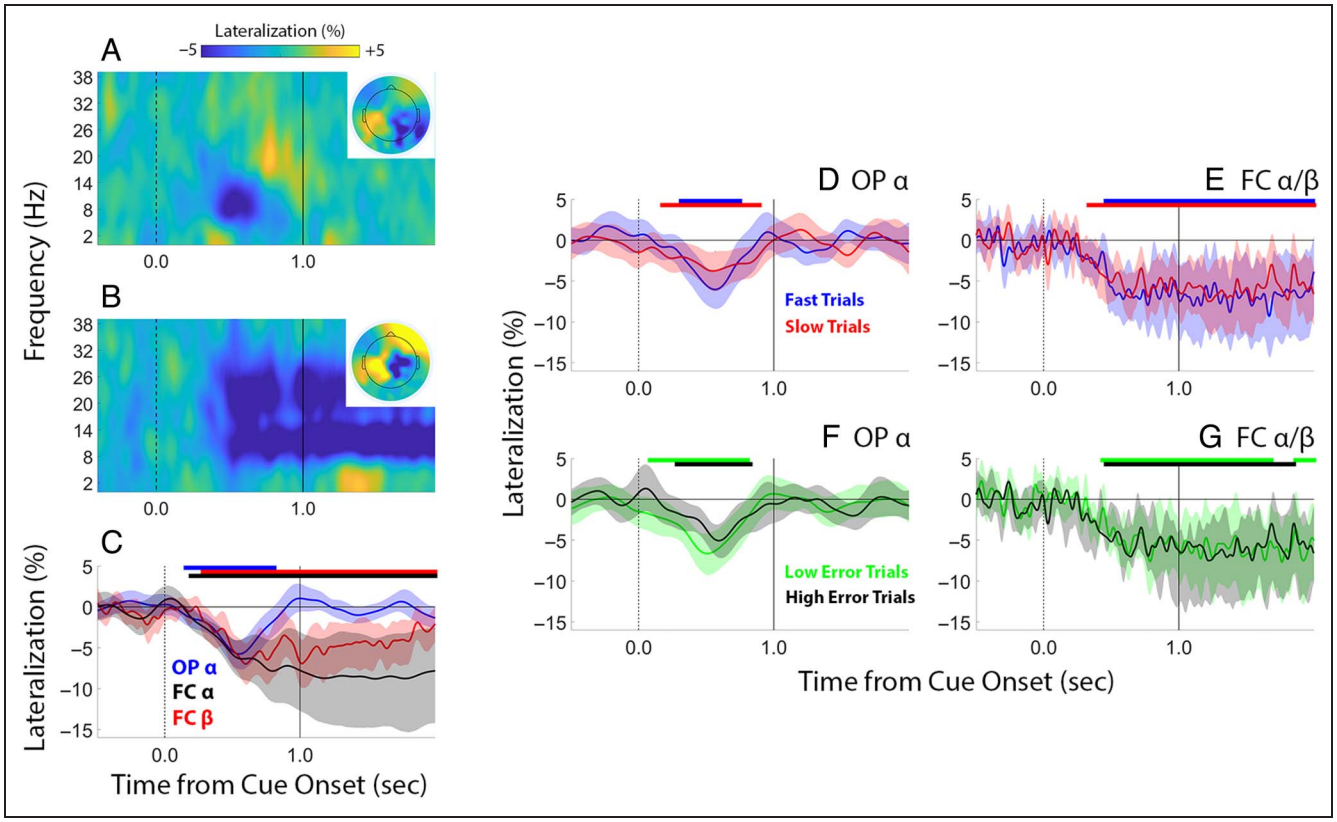

Figure 2. EEG signatures of spatial and motor selection. (A) Spectral signals that co-varied with stimulus location (i.e., left vs. right visual field). Lateralization estimates were computed from occipitoparietal electrode site pairs O1/2, PO3/4, and PO7/8. The inset shows the distribution of lateralized alpha power (8–13 Hz) across all scalp electrodes, averaged over a period spanning 0.5–1.0 sec after cue onset. (B) Spectral signals that co-varied with response demands (i.e., left vs. right hand). Lateralization estimates were computed from frontocentral site pairs C1/2 and C3/4. The inset shows the distribution of lateralized alpha/beta power (8–13 Hz) across all scalp electrodes, averaged over a period spanning 0.5–1.0 sec after cue onset. (C) Direct comparisons of occipitoparietal alpha (OP α), frontocentral mu-alpha (FC α), and frontocentral mu-beta (FC β) power. Trial-wise differences in occipitoparietal (OP) α power (8–13 Hz) and frontocentral (FC) α/β power (8–30 Hz) did not predict trial-wise differences in response onset (D and E, respectively) or memory precision (F and G, respectively). Horizontal bars at the top of C–G depict epochs where EEG lateralization was significantly less than zero (cluster-corrected permutation tests). Shaded regions depict the 95% confidence interval of the mean.

following probe onset; Figure 2D–E) nor participants' memory performance (i.e., recall error; Figure 2F–G).

Oculomotor Signatures of Spatial and Motor Selection

A handful of recent studies have reported small but robust biases in gaze (e.g., van Ede, Checkroud, & Nobre, 2019) and head direction (Thom, Nobre, van Ede, & Draschkow, 2023) toward the physical location of a retrospectively cued stimulus held in WM, even when location is task irrelevant. Importantly, these gaze biases appear to be at least partially independent of EEG signals associated with covert spatial attention (e.g., Liu, Nobre, & van Ede, 2022; see also Yu et al., 2022, for an analogous finding in nonhuman primates). Motivated by these findings, we sought to compare the amplitudes and latencies of EEG and oculomotor signals of covert spatial attention. We obtained high-quality horizontal gaze position recordings for 18 of the 31 participants who completed this study (EEG and gaze position data were acquired concurrently; see Methods). Following earlier work (van Ede, Checkroud, & Nobre, 2019), we converted participants' horizontal gaze position measurements into a normalized metric where 0% represents perfect fixation and −100% represents foveating the center of the cued stimulus (we used a negative scale to facilitate visual comparisons with negative-going changes in EEG lateralization; Figure 3). Consistent with earlier findings (e.g., van Ede, Checkroud, & Nobre, 2019), we observed a small (~3%) but robust bias in horizontal gaze position toward the cued stimulus beginning 282 msec after cue onset (Figure 3A).

In principle, gaze biases present during visuomotor selection could reflect the attentional selection of the cued stimulus' position, the cued response hand, or some mixture of both. Although our study was not designed to disentangle these possibilities, we reasoned that if gaze position biases are jointly influenced by spatial and motor selection, then they should be larger during congruent trials where the selection of spatial and motor information is aligned (e.g., a left visual field stimulus requiring a left hand response) compared with incongruent trials where the selection of spatial and motor information is misaligned (e.g., a left visual field stimulus requiring a right-

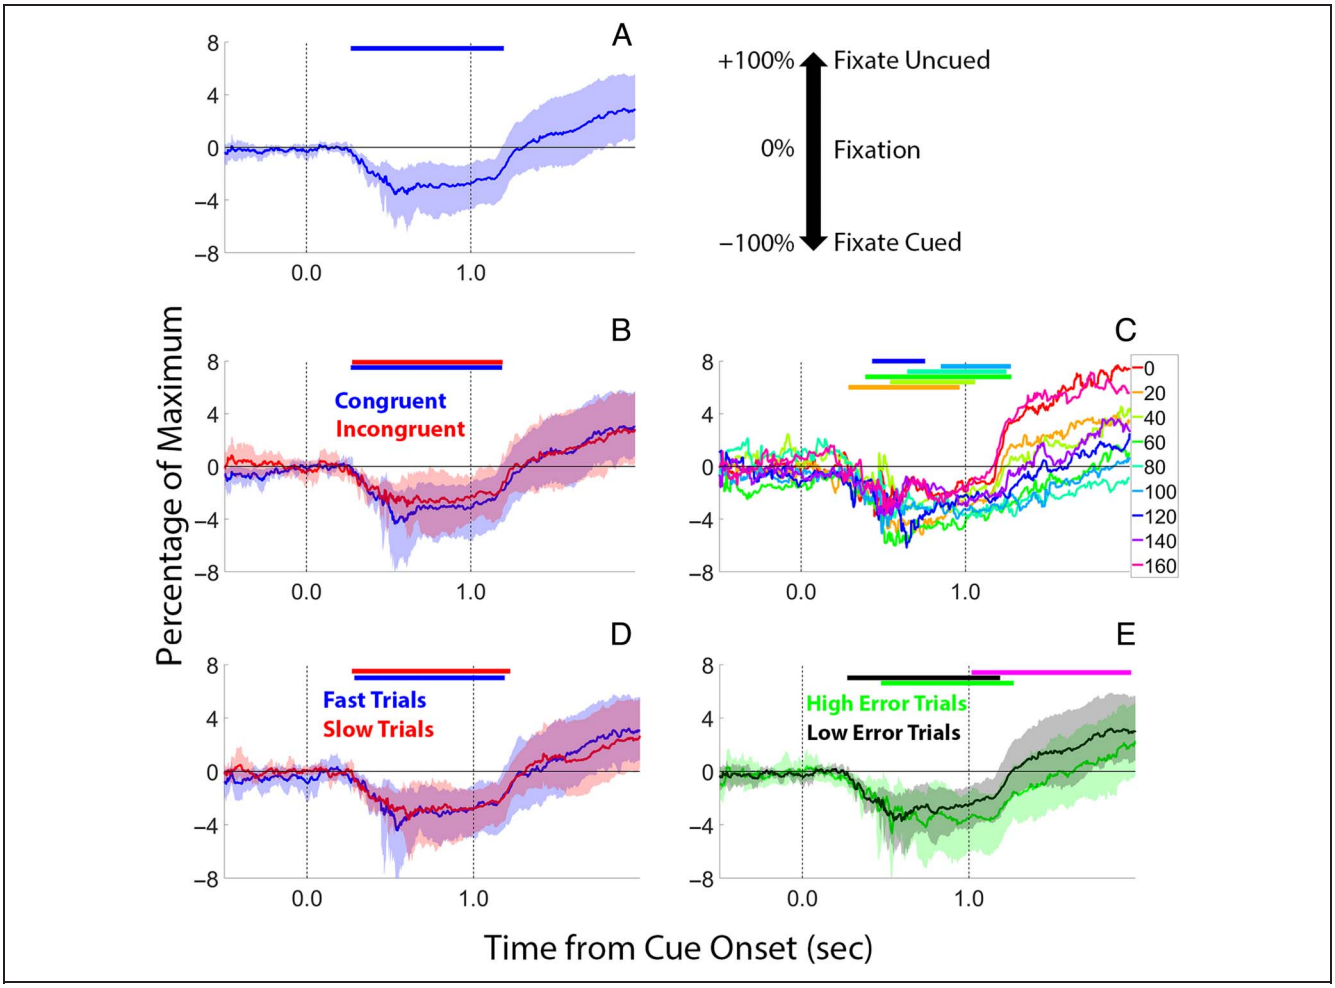

Figure 3. Horizontal gaze biases associated with spatial, motor, and feature selection. (A) Even though the parameters of our task rendered spatial information irrelevant (i.e., participants could solve the task by remembering color orientation bindings regardless of their positions during encoding), we observed a small but robust horizontal bias toward the position of the retrospectively cued stimulus. Gaze biases are expressed as a percentage of the maximum possible bias, with −100% corresponding to fixating the retrospectively cued grating and +100% corresponding to fixating the noncued grating. Shifts toward the cued grating are plotted downward to facilitate comparisons with the negative-going EEG lateralization estimates plotted in Figure 2. (B) Gaze biases were equivalent in magnitude and duration during congruent trials where spatial and motor selection were aligned (e.g., a left visual field stimulus requiring a left-hand response) and incongruent trials where spatial and motor selection were misaligned (e.g., a left visual field stimulus requiring a right-hand response), suggesting that gaze biases were driven exclusively by the cued stimulus position. (C) Gaze biases during the cue-to-probe interval were unaffected by stimulus orientation. Trial-wise differences in gaze bias did not predict variability in response onset (D) during either the cue-to-probe or recall periods. Trial-wise differences in gaze bias did not predict variability in recall error (E) during the cue-to-probe interval but did predict recall error during the recall period, with low-error trials associated with reliably smaller gaze position biases during the recall period compared with high-error trials. Horizontal bars at the top of each plot depict epochs where gaze position was significantly less than zero (i.e., biased toward the location of the retrospectively cued disc; cluster-corrected permutation tests). Shaded regions depict the 95% confidence interval of the mean; shaded regions were omitted from C to improve readability. Vertical lines at times 0.0 and 1.0 depict the onset of the cue and probe displays, respectively.

hand response). However, direct comparisons of gaze biases during congruent and incongruent trials revealed no differences in either the magnitude nor the duration of horizontal gaze biases (Figure 3B), suggesting that these biases are driven primarily by the spatial position of the to-be-recalled stimulus and not the preparation of a manual response.

Finally, we considered the possibility that gaze position biases were jointly influenced by stimulus position and orientation. We tested this possibility by replotting and analyzing horizontal gaze bias as a function of stimulus position (left vs. right hemifield) and stimulus orientation (0–160° in 20° increments; Figure 3C). This analysis revealed a significant modulation of gaze position biases by stimulus orientation during the later portion of the recall period (250–500 msec after probe onset; one-way repeated-measures ANOVA with Stimulus Orientation as the sole factor; $F(8, 136) = 3.995$, $p = .0003$, $\eta^2 = .190$), but not during the early portion of the recall period (0–250 msec after probe onset; $F(8, 136) = 1.412$, $p = .197$, $\eta^2 = .077$) nor during the cue-to-probe interval (e.g., 500–1000 msec after retrocue onset; $F(8, 136) = 0.778$, $p = .623$, $\eta^2 = .044$). Like the EEG data, gaze biases during the cue-to-probe interval did not predict participants'

response speeds (i.e., time to first keypress following probe onset; Figure 3D) nor participants' memory performance (i.e., recall error; Figure 3E).

EEG Signatures of Feature Selection

Several lines of evidence indicate that the brain automatically retains spatial information in WM: First, location-specific representations can be reconstructed from scalp EEG activity during WM tasks where location is irrelevant (Foster et al., 2017). Second, human gaze (e.g., van Ede, Checkroud, & Nobre, 2019) and head position (Thom et al., 2023) are biased toward the location of a behaviorally prioritized item stored in WM, also in tasks where location is irrelevant. Third, changes in gaze position or shifts of covert spatial attention during WM can disrupt or introduce systematic biases in memory for visual features (e.g., color; Golomb, Kupitz, & Thiemann, 2014), and memory-guided comparisons for successively presented visual stimuli are impaired when they are rendered in different spatial positions (e.g., Hollingworth, 2007). Automatic storage of visuospatial information in WM may be adaptive, enabling the segregation of mental representations in retinotopic or spatiotopic coordinate frames during memory encoding, storage, and retrieval (Groen et al., 2022). Our data support this perspective: Even though our experimental task (Figure 1) did not oblige participants to retain spatial information in WM, we nevertheless observed robust electrophysiological (Figure 2) and oculomotor (Figure 3) evidence for the selection of spatial information following the appearance of a retrospective cue.

We reasoned that if visuospatial information plays an anchoring function during WM storage and WM-guided behavior, then the selection of spatial information might also prompt the selection of task-relevant feature information, in this case, orientation. We examined this possibility by testing whether it was possible to decode the orientations of the cue-matching and -nonmatching gratings during the cue-to-probe interval when evidence for spatial selection was highest (e.g., Figure 2C). Surprisingly, although we observed clear EEG and oculomotor evidence for spatial (and motor) selection during the cue-to-probe interval, above-chance decoding of the retrospectively cued orientation only emerged during the recall period (Figure 4). This general pattern held when we decoded stimulus orientation from the same occipitoparietal electrodes used to quantify spatial selection (O1/2, PO3/4, and PO7/8; Figures 2C and 4A), the same frontocentral electrode sites used to quantify motor selection (FC1/2; FC3/4; Figures 2C 4B) or all 62 scalp electrodes (Figure 4C).

We considered several explanations for the absence of above-chance orientation decoding during the cue-to-probe interval, as well as the presence of above-chance orientation decoding during the recall interval. First, we considered the possibility that above-chance decoding of orientation during the recall period reflects participants' adjusting the orientation of the probe grating to match the orientation of the cued grating rather than the selection of orientation information stored in WM per se. We tested this possibility by comparing the onset of above-chance orientation decoding performance with the timing of participants' responses (i.e., the amount of time that elapsed between the onset of the probe display and the time at which participants began to adjust the probe stimulus, i.e., 451 ± 14.9 msec). The onset of above-chance decoding performance occurred significantly earlier than the average response onset for the data shown in Figure 4A (mean above-chance decoding onset relative to the appearance of the probe display = 182 msec;

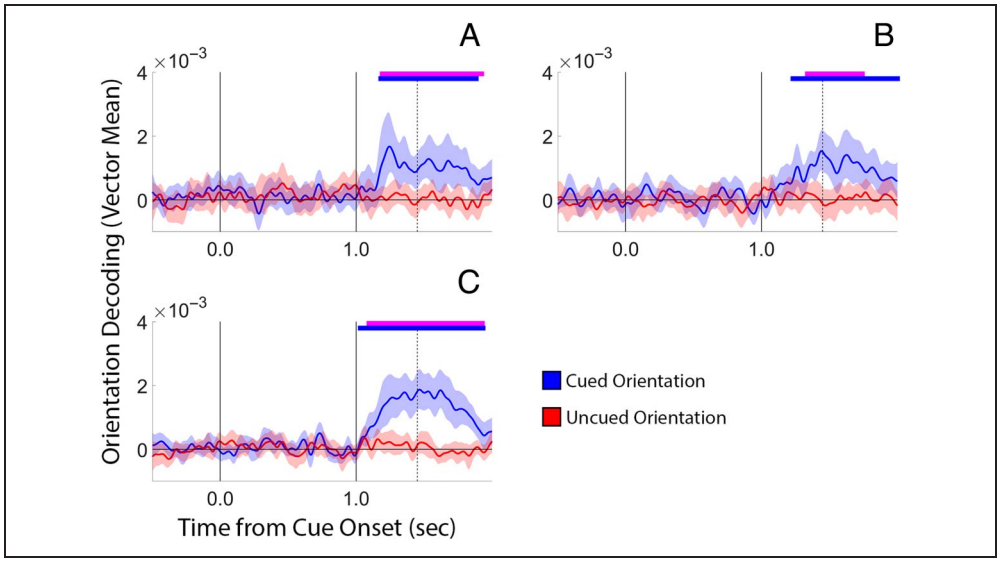

Figure 4. Orientation decoding time-locked to retrocue onset. Orientation decoding was computed using data from occipitoparietal electrode site pairs O1/2, PO3/4, and PO7/8 (A); frontocentral site pairs FC1/2 and FC3/4 (B); or all 62 scalp electrodes (C). Horizontal bars at the top of each plot mark epochs where decoding performance was significantly greater than chance; pink horizontal bars mark epochs where cued orientation decoding performance was significantly greater than uncued orientation decoding performance. The solid vertical lines at times 0.0 and 1.0 depict the onset of the retrocue and probe displays, respectively. The dashed vertical line at time 1.451 sec marks participants' average response onset, that is, the point at which they first began to rotate the probe grating. Shaded regions in each plot depict the 95% confidence interval of the mean.

p = .0060; bootstrap test against a distribution with a mean of 451 msec), Figure 4B (mean above-chance decoding onset = 234 msec; p = .0004), and Figure 4C (mean above-chance decoding onset = 30 msec; p = .0001). Thus, above-chance decoding of orientation cannot be explained by participants adjusting the probe grating to match that of the cued grating.

Second, we considered the possibility that above-chance orientation decoding during the early portion of the response period (i.e., before response onset) was driven by the mere appearance of an oriented stimulus. However, note that the probe grating was always rendered with a vertical orientation, and thus, across trials the initial orientation of the probe grating was randomized with respect to the orientation of the remembered grating (drawn from the set [0°, 20°, 40°, 60°, 80°, 100°, 120°, 140°, 160°]). This argues against the possibility that above-chance decoding during the initial part of the recall period was driven by the mere appearance of an oriented stimulus (we return to this point in the discussion).

Third, we considered the possibility that the absence of above-chance orientation decoding during the cue-to-probe interval—when participants were free to select the cue-matching orientation for upcoming report and when EEG and oculomotor evidence for spatial and motor selection was evident (e.g., Figures 2 and 4)—was due to a lack of analytic sensitivity. For example, one recent study reported that the orientation of a remembered stimulus could be decoded from alpha-band but not broadband EEG activity during WM storage (Barbosa, Lozano-Soldevilla, & Compte, 2021). We therefore repeated our decoding analysis after applying an 8–13 Hz bandpass filter to the EEG data. The results of this analysis were remarkably similar to those summarized in Figure 4, with robust above-chance decoding of orientation observed during the recall period but not the cue-to-probe period (Figure 5A–C).

Finally, we considered the possibility that the absence of above-chance orientation decoding during the cue-to-probe interval was caused by an analytic artifact. For example, one recent study demonstrated that high-pass filter cutoffs over 0.1 Hz can systematically bias the time(s) at which feature information can be decoded from broadband EEG data (van Driel, Olivers, & Fahrenfort, 2021). We tested this possibility by repreprocessing our data with a 0.1-Hz high-pass filter cutoff and applying our decoding analysis to the resulting data. The results of this analysis were remarkably similar to the data summarized in Figure 4 (Figure 5D–F), suggesting that the absence of above-chance decoding performance during the cue-to-probe interval cannot be explained by a filtering artifact.

Eye Movement Control Analyses

Recent studies suggest that subtle differences in gaze position can contribute to feature decoding performance during WM storage and readout (e.g., Quax et al., 2019; Mostert et al., 2018). We therefore sought to understand whether and how gaze position may have influenced EEG decoding in the current study. As a first approximation, we tested whether horizontal gaze position biases were modulated by stimulus orientation during the cue-to-

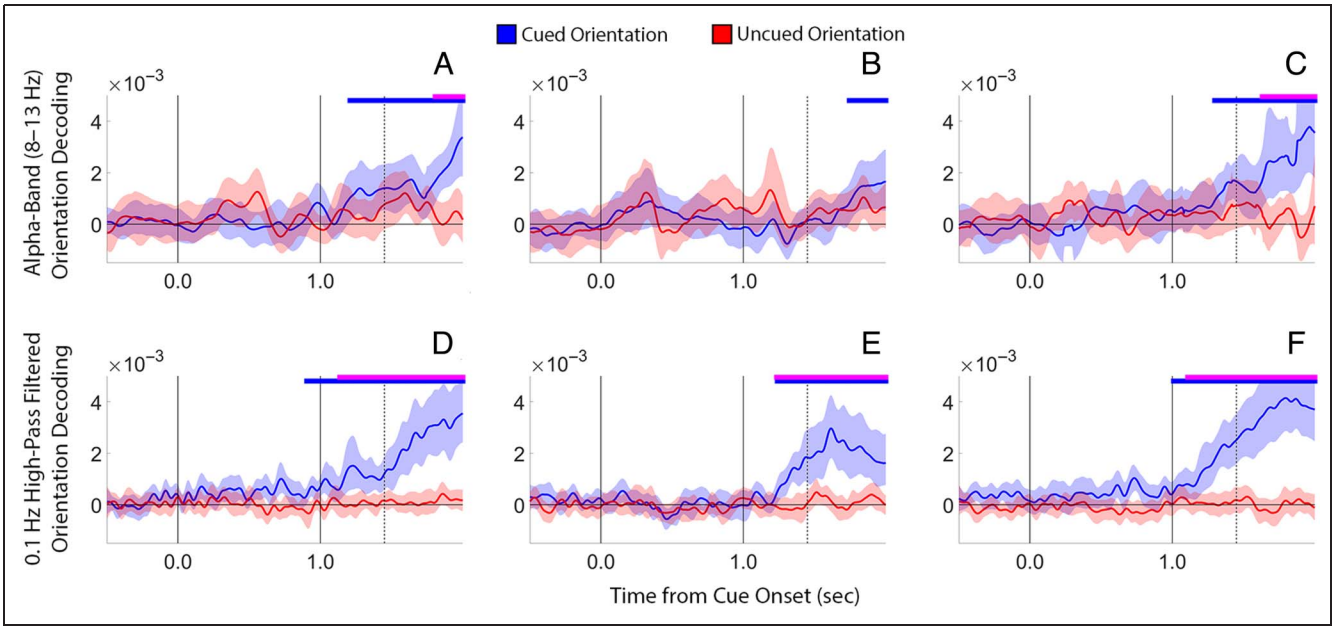

Figure 5. Supplementary orientation decoding analyses. (A–C) Orientation decoding computed from alpha-band filtered EEG data over occipitoparietal electrodes (A), frontocentral electrodes (B), and all 62 scalp electrodes (C). (D–F) Orientation decoding performance computed from modestly (0.1 Hz) high-pass filtered EEG data over occipitoparietal electrodes (D), frontocentral electrodes (E), and all 62 scalp electrodes (F). All conventions are identical to Figure 4.

probe interval (i.e., −250 to −1 msec before probe onset) and during the recall period (i.e., 1–250 msec after probe onset). We averaged participants' horizontal gaze position estimates as a function of stimulus orientation and subjected the resulting values to a one-way repeated-measures ANOVA with Stimulus Orientation as the sole factor. This analysis did not reveal an effect of Orientation during the cue-to-probe interval (when it was not possible to decode stimulus orientation from EEG signals; $F(8, 136) = 0.778$, $p = .623$, $\eta^2 = .044$) or the recall period (when it was possible to decode stimulus orientation from EEG signals; $F(8, 136) = 1.412$, $p = .197$, $\eta^2 = .077$). However, additional analyses did reveal a robust effect of stimulus orientation on horizontal gaze position during later portions of the recall period (e.g., 250–500 msec after probe onset; $F(8, 136) = 3.995$, $p = .0003$, $\eta^2 = .190$). Thus, we undertook additional analyses to understand how differences in gaze position may have influenced EEG decoding performance.

To obtain a more direct estimate of the relationship between gaze position and stimulus orientation, we attempted to decode the latter from the former. Specifically, we used the same parametric distance-based decoding procedure used to decode stimulus orientation from EEG signals to decode stimulus orientation from trial- and time-wise records of gaze position in $[x, y]$ screen coordinates. Because of a hardware malfunction, vertical gaze position data from one participant were lost during recording. Thus, we restricted our analysis to the remaining 17 participants with concurrent and robust EEG and $[x, y]$ gaze position recordings. This analysis revealed significant above-chance gaze position-based decoding of cued orientation during the recall period (Figure 6A). However, an examination of individual subject data indicated that above-chance decoding was driven primarily by three extreme participants (Figure 6B, red traces). Following earlier work (Printzlau, Myers, Manohar, & Stokes, 2022), we examined the influence of these participants on gaze position-based orientation decoding performance by computing average decoding performance during the probe-to-response period (i.e., 0–451 msec after probe onset, during which time it was possible to decode stimulus orientation from EEG data; Figure 4) while including and excluding suspected outliers from the analysis. As expected, gaze position-based orientation decoding performance fell to chance levels when the three extreme participants highlighted in Figure 6B were excluded from the analysis (Figure 6C; two-tailed t test against chance, i.e., 0: $t(13) = 1.96$, $p = .072$). Finally, we confirmed that EEG-based decoding performance during the probe-to-response interval remained well above chance levels even when participants with extremely high gaze position-based decoding performance were excluded from the analysis (Figure 6D), indicating that gaze position biases

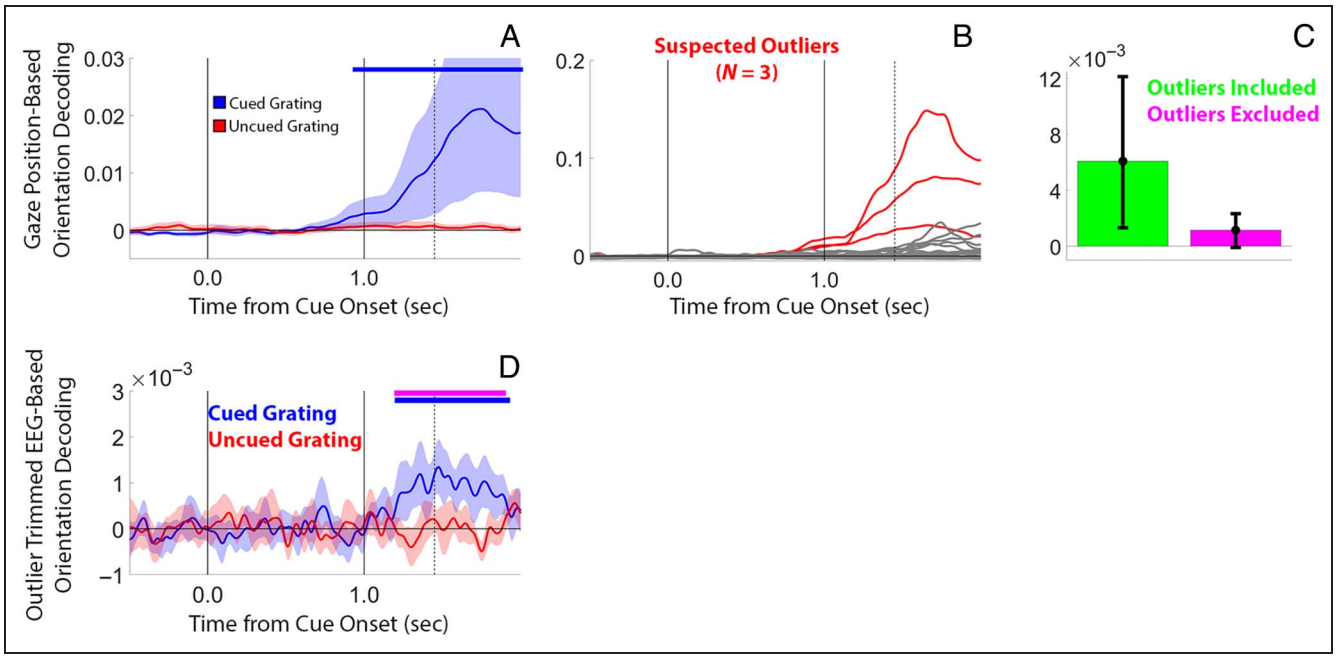

Figure 6. Eye movements contribute to but are not fully responsible for EEG orientation decoding performance. (A) Average orientation decoding performance computed using gaze position data (x, y coordinates). (B) Participant-level orientation decoding performance computed using gaze position data. Suspected outliers are highlighted in red. The solid vertical lines at times 0.0 and 1.0 depict the onset of the retrocue and probe displays, respectively. The vertical dashed line at time 1.451 sec. Depicts the average response onset across participants. (C) Average orientation decoding performance computed from gaze position data over a period spanning 0–451 msec after probe onset while including versus excluding the suspected outliers shown in B. (D) Orientation decoding performance computed from EEG data from 14 of 17 participants with concurrent EEG and gaze position recordings (i.e., excluding the three possible outliers shown in B. Conventions are identical to A and B. Shaded regions and error bars in each plot depict the 95% confidence interval of the mean.

were not solely responsible for above-chance EEG-based orientation decoding during this period.

DISCUSSION

Human memory systems exist so that prior experiences can guide immediate and future behaviors. In recognition of this fact, recent conceptualizations of short-term memory and/or WM have begun to emphasize the action-oriented nature of this system. Recent findings demonstrate that WM and action planning interact in a bidirectional manner: The contents of WM can influence movement planning and execution (e.g., see van Ede & Nobre, 2023, for a recent review), and incidental movements during storage can influence the contents of WM (Heuer et al., 2020; Heuer & Schubo, 2017). Other recent findings suggest that the human brain can store multiple stimulus–response mappings in WM and that behaviorally relevant WM contents can be selected in parallel with required actions (van Ede, Chekroud, Stokes, et al., 2019). Although this result seems to challenge classic "think-then-act" conceptualizations of human memory performance (e.g., Donders, 1969; Sternberg, 1969), in our view the evidence supporting this challenge is limited. Here, we re-examined the chronometry of spatial, feature, and motor selection during WM-guided behavior during a task optimized to disentangle neural signals associated with these processes. We found clear evidence for a temporal dissociation between feature and visuomotor selection: Although EEG signatures linked to spatial and motor selection appeared coincidentally shortly after the appearance of an informative recall cue (Figure 2), EEG signatures linked to feature selection appeared significantly later (Figure 4). Our findings reveal an important distinction between spatial and nonspatial selection during WM-guided behavior and clarify how these selection mechanisms align with the selection of item-specific motor plans. More generally, our findings challenge conceptualizations of WM-guided behavior based exclusively on serial versus parallel visuomotor processing.

Intuitively, the visual features of a task-relevant object could be selected automatically along with its location, complementing demonstrations of object-based selection in WM where selecting one nonspatial attribute of an object (e.g., color) leads to the selection of other nonspatial attribute of the same object (e.g., orientation; Printzlau et al., 2022). Alternatively, given the apparent primacy of spatial position in organizing visual perception and WM (e.g., Groen et al., 2022), the selection of feature information stored in WM might be functionally and/or temporally decoupled from the selection of spatial information stored in WM. Our data support this view. Specifically, although we observed clear EEG and eye position evidence for spatial selection during the cue-to-probe interval (Figures 2 and 3), evidence for feature selection was limited to the recall period (Figure 4). Thus, our data reveal the existence of temporally distinct mechanisms for the selection of spatial and nonspatial visual information during WM-guided behavior.

In the current study, we found that feature selection (measured via orientation decoding) lagged spatial and motor selection (measured via lateralized EEG signals) by several hundred milliseconds. One possibility is that participants selected spatial and motor information during the cue-to-probe period but waited to select orientation information until the response period. Although we cannot fully exclude this possibility, we think that it is unlikely. Holding multiple items in WM is known to reduce the item-specific information content of multivariate EEG and fMRI signals (e.g., Ester et al., 2018; Sprague, Ester, & Serences, 2014). However, this reduction can be partially reversed by prioritizing one item stored in memory (e.g., Ester & Pytel, 2023; Ester et al., 2018; Sprague, Ester, & Serences, 2016). Because our task required participants to recall orientations with the highest possible precision, the optimal strategy is to select the task-relevant orientation as soon as it is possible to do so. That participants did not or could not do so points to a temporal dissociation between spatial, feature, and motor selection.

One possibility is that the selection of feature information from WM can only proceed after the appearance of a trailing sensory stimulus. Activity-silent models of WM propose that item-specific information is stored via short-term changes in synaptic weights in a way that permits trailing sensory inputs to reactivate neural firing patterns seen during WM encoding (e.g., Wolff et al., 2017; Stokes, 2015; Mongillo, Barak, & Tsodyks, 2008). It is unclear, however, whether short-term changes in synaptic weights are restricted to the retinotopic locations where remembered stimuli appeared. In prior studies supporting activity-silent WM, trailing stimuli used to reactivate latent WM representations always appeared in the same retinotopic locations as remembered stimuli (e.g., Wolff et al., 2017); in the current study the remembered stimulus and the probe stimulus were always rendered in different retinotopic locations (i.e., laterally vs. foveally, respectively; Figure 1). One way to reconcile this apparent discrepancy is to propose that activity-silent WM is spatially global, that is, the changes in synaptic weights responsible for storing information extend beyond the retinotopic locations where remembered stimuli appeared. This arrangement would dovetail with other observations of spatially global WM, where feature information can be decoded from neural populations retinotopically mapped versus not retinotopically mapped to the location of a remembered stimulus (e.g., Ester, Serences, & Awh, 2009). Other hypotheses could explain the emergence of above-chance orientation decoding during the recall period. For example, perhaps feature selection is in fact temporally dissociated from spatial and motor selection and the probe stimulus acts as a "second cue" that warns participants about upcoming response demands and thus prompts a voluntary selection of the task-relevant orientation. Although these alternatives await further scrutiny, we

note that both would be inconsistent with a purely parallel model where the selection of spatial information automatically prompts the selection of accompanying feature information.

Individual variability in WM is strongly correlated with general cognitive function (e.g., IQ; Cowan, 2001), and studies have shown that participants with high WM function efficiently gate access to this system (e.g., Feldmann-Wüstefeld & Vogel, 2019). Given that WM is a fundamentally action-oriented system, the ability to efficiently select content already in WM and use this information for behavior is likely to play an equally important role in cognitive performance. Here, we used EEG and eye position data to track the selection of spatial, feature, and motor information during WM-guided behavior. Our data reveal the presence of multiple temporally dissociable selection mechanisms for WM-guided behavior, for example, with spatial and motor selection occurring before the read-out of feature information from WM. More generally, our findings refute conceptualizations of WM-guided behavior based purely on serial versus parallel visuomotor processing.

Reprint request should be sent to Edward Ester, University of Nevada, Reno, 1664 N. Virginia St., MC 0296, Reno, NV 89557, or via e-mail: eester@unr.edu.

Data Availability Statement

Stimulus presentation software, preprocessed data files, and analytic software sufficient to reproduce all manuscript figures and reported statistical values are publicly available on OSF: https://osf.io/dm73j/.

Author Contributions

Edward Ester: Conceptualization; Formal analysis; Funding acquisition; Investigation; Methodology; Project administration; Software; Supervision; Visualization; Writing—Original draft; Writing—Review & editing. Rachel Weese: Conceptualization; Formal analysis; Investigation; Methodology; Software; Writing—Review & editing.

Funding Information

National Science Foundation (https://dx.doi.org/10.13039/100000001), grant number: 2050833.

Diversity on Citation Practices

A retrospective analysis of the citations in every paper published in this journal from 2010–2020 has revealed a persistent pattern of gender imbalance: Although the proportions of authorship teams (categorized by estimated gender identification of first-author/last-author) publishing in *Journal of Cognitive Neuroscience (JoCN)* during this time period were M(an)/M = .408 W(oman)/M = .335 M/W = .108 W/W = .149; the comparable proportions for the papers that these authorship teams cited were M/M = .579 W/M = .243 M/W = .102 W/W = .076 (Fulvio et al., *JoCN*, 33:1, pp. 3–7). Consequently, *JoCN* encourages all authors to consider gender balance explicitly when selecting which papers to cite, and gives them the opportunity to report their paper's gender citation balance. The authors of this paper report its proportions of citations by gender category to be: M/M = .674; W/M = .116; M/W = .163; W/W = .047.

REFERENCES

Barbosa, J., Lozano-Soldevilla, D., & Compte, A. (2021). Pinging the brain with visual impulses reveals electrically active, not activity-silent, working memories. *PLoS Biology*, *19*, e3001436. https://doi.org/10.1371/journal.pbio.3001436, PubMed: 34673775

Boettcher, S. E. P., Gresch, D., Nobre, A. C., & van Ede, F. (2021). Output planning at the input stage in visual working memory. *Science Advances*, *7*, eabe8212. https://doi.org/10.1126/sciadv.abe8212, PubMed: 33762341

Chang, C.-Y., Hsu, S.-H., Pion-Tonachini, L., & Jung, T.-P. (2020). Evaluation of Artifact Subspace Reconstruction for automatic artifact components removal in multi-channel EEG recordings. *IEEE Transactions on Biomedical Engineering*, *67*, 1114–1121. https://doi.org/10.1109/TBME.2019.2930186, PubMed: 31329105

Chatham, C. H., Frank, M. J., & Badre, D. (2014). Corticostriatal output gating during selection from working memory. *Neuron*, *81*, 930–942. https://doi.org/10.1016/j.neuron.2014.01.002, PubMed: 24559680

Cowan, N. (2001). The magical number 4 in short-term memory: A reconsideration of mental storage capacity. *Behavioral and Brain Sciences*, *24*, 87–114. https://doi.org/10.1017/S0140525X01003922, PubMed: 11515286

Delorme, A., & Makeig, S. (2004). EEGLAB: An open-source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, *134*, 9–21. https://doi.org/10.1016/j.jneumeth.2003.10.009, PubMed: 15102499

Donders, F. C. (1969). On the speed to mental processes. In W. G. Koster (Ed.), *Attention and performance II* (pp. 412–431). Amsterdam: North-Holland. https://doi.org/10.1016/0001-6918(69)90065-1

Ester, E. F., Nouri, A., & Rodriguez, L. (2018). Retrospective cues mitigate information loss in human cortex during working memory storage. *Journal of Neuroscience*, *38*, 8538–8548. https://doi.org/10.1523/JNEUROSCI.1566-18.2018, PubMed: 30126971

Ester, E. F., & Pytel, P. (2023). Changes in behavioral priority influence the accessibility of working memory content. *Neuroimage*, *272*, 120055. https://doi.org/10.1016/j.neuroimage.2023.120055, PubMed: 37001833

Ester, E. F., Serences, J. T., & Awh, E. (2009). Spatially global representations in human primary visual cortex during working memory maintenance. *Journal of Neuroscience*, *29*, 15258–15265. https://doi.org/10.1523/JNEUROSCI.4388-09.2009, PubMed: 19955378

Feldmann-Wüstefeld, T., & Vogel, E. K. (2019). Neural evidence for the contribution of active suppression during working memory filtering. *Cerebral Cortex*, *29*,

529–543. https://doi.org/10.1093/cercor/bhx336, PubMed: 29365078
Foster, J. J., Bsales, E. M., Jaffe, R. J., & Awh, E. (2017). Alpha-band activity reveals spontaneous representations of spatial position in visual working memory. *Current Biology*, *27*, 3216–3223. https://doi.org/10.1016/j.cub.2017.09.031, PubMed: 29033335
Galero-Salas, Y., Han, S., Sych, Y., Voigt, F. F., Laurenczy, B., Gilad, A., et al. (2021). Sensory and behavioral components of neocortical signal flow in discrimination tasks with short-term memory. *Neuron*, *109*, 135–148. https://doi.org/10.1016/j.neuron.2020.10.017, PubMed: 33159842
Golomb, J. D., Kupitz, C. N., & Thiemann, C. T. (2014). The influence of object location on identity: A "spatial congruency bias.". *Journal of Experimental Psychology: General*, *143*, 2262–2278. https://doi.org/10.1037/xge0000017, PubMed: 25222263
González-García, C., Formica, S., Liefooghe, B., & Brass, M. (2020). Attentional prioritization reconfigures novel instructions into action-oriented task sets. *Cognition*, *194*, 104059. https://doi.org/10.1016/j.cognition.2019.104059, PubMed: 31514104
Groen, I. I. A., Dekker, T. M., Knapen, T., & Silson, E. H. (2022). Visuospatial coding as ubiquitous scaffolding for human cognition. *Trends in Cognitive Sciences*, *26*, 81–96. https://doi.org/10.1016/j.tics.2021.10.011, PubMed: 34799253
Heuer, A., Ohl, S., & Rolfs, M. (2020). Memory for action: A functional view of selection in visual working memory. *Visual Cognition*, *28*, 388–400. https://doi.org/10.1080/13506285.2020.1764156
Heuer, A., & Schubo, A. (2017). Selective weighting of action-related feature dimensions in visual working memory. *Psychonomic Bulletin & Review*, *24*, 1129–1134. https://doi.org/10.3758/s13423-016-1209-0, PubMed: 27896631
Hollingworth, A. (2007). Object-position binding in visual memory for natural scenes and object arrays. *Journal of Experimental Psychology: Human Perception and Performance*, *33*, 31–47. https://doi.org/10.1037/0096-1523.33.1.31, PubMed: 17311477
Kaiser, J., Birbaumer, N., & Lutzenberger, W. (2001). Event-related beta desynchronization indicates timing of response selection in a delayed-response paradigm in humans. *Neuroscience Letters*, *312*, 149–152. https://doi.org/10.1016/S0304-3940(01)02217-0, PubMed: 11602332
Keefe, J. M., & Störmer, V. S. (2021). Lateralized alpha activity and slow potential shifts over visual cortex track the time course of both endogenous and exogenous orienting of attention. *Neuroimage*, *225*, 117495. https://doi.org/10.1016/j.neuroimage.2020.117495, PubMed: 33184032
Kleiner, M., Brainard, D., & Pelli, D. (2007). What's new in Psychtoolbox-3? *Perception*, *36*, 1–16. https://hdl.handle.net/21.11116/0000-0004-44C8-E
Klimesch, W. (2012). Alpha-band oscillations, attention, and controlled access to stored information. *Trends in Cognitive Sciences*, *16*, 606–617. https://doi.org/10.1016/j.tics.2012.10.007, PubMed: 23141428
Kopell, N., Whittington, M. A., & Kramer, M. A. (2011). Neuronal assembly dynamics in the beta1 frequency range permits short-term memory. *Proceedings of the National Academy of Sciences, U.S.A.*, *108*, 3779–3784. https://doi.org/10.1073/pnas.1019676108, PubMed: 21321198
LaRocque, J. J., Lewis-Peacock, J. A., Drysdale, A. T., Oberauer, K., & Postle, B. R. (2013). Decoding attended information in short-term memory: An EEG study. *Journal of Cognitive Neuroscience*, *25*, 127–142. https://doi.org/10.1162/jocn_a_00305, PubMed: 23198894
Liu, B., Nobre, A. C., & van Ede, F. (2022). Functional but not obligatory link between microsaccades and neural modulation by covert spatial attention. *Nature Communications*, *13*, 3503. https://doi.org/10.1038/s41467-022-31217-3, PubMed: 35715471
Maris, E., & Oostenveld, R. (2007). Nonparamteric statistical testing of EEG- and MEG-data. *Journal of Neuroscience Methods*, *164*, 177–190. https://doi.org/10.1016/j.jneumeth.2007.03.024, PubMed: 17517438
Meyer, D. E., Osman, A. M., Irwin, D. E., & Yantis, S. (1988). Modern mental chronometry. *Biological Psychology*, *26*, 3–67. https://doi.org/10.1016/0301-0511(88)90013-0, PubMed: 3061480
Miller, E. K., Lundqvist, M., & Bastos, A. M. (2018). Working Memory 2.0. *Neuron*, *100*, 463–475. https://doi.org/10.1016/j.neuron.2018.09.023, PubMed: 30359609
Mongillo, G., Barak, O., & Tsodyks, M. (2008). Synaptic theory of working memory. *Science*, *319*, 1543–1546. https://doi.org/10.1126/science.1150769, PubMed: 18339943
Mostert, P., Albers, A. M., Brinkman, L., Todorova, L., Kok, P., & de Lange, F. P. (2018). Eye movement-related confounds in neural decoding of visual working memory representations. *eNeuro*, *5*, ENEURO.0401-17.2018. https://doi.org/10.1523/ENEURO.0401-17.2018, PubMed: 30310862
Ohl, S., & Rolfs, M. (2020). Bold moves: Inevitable saccadic selection in visual short-term memory. *Journal of Vision*, *20*, 11. https://doi.org/10.1167/jov.20.2.11, PubMed: 32106297
Olivers, C. N. L., & Roelfsema, P. R. (2020). Attention for action in visual working memory. *Cortex*, *131*, 179–194. https://doi.org/10.1016/j.cortex.2020.07.011, PubMed: 32892152
O'Reilly, R. C., & Frank, M. J. (2006). Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Computation*, *18*, 283–328. https://doi.org/10.1162/089976606775093909, PubMed: 16378516
Perrin, F., Pernier, J., Bertrand, O., & Echallier, J. F. (1989). Spherical splines for scalp potential and current density mapping. *Electroencephalography and Clinical Neurophysiology*, *72*, 184–187. https://doi.org/10.1016/0013-4694(89)90180-6, PubMed: 2464490
Pho, G. N., Goard, M. J., Woodson, J., Crawford, B., & Sur, M. (2018). Task-dependent representations of stimulus and choice in mouse parietal cortex. *Nature Communications*, *9*, 2596. https://doi.org/10.1038/s41467-018-05012-y, PubMed: 29968709
Printzlau, F. A. B., Myers, N. E., Manohar, S. G., & Stokes, M. G. (2022). Neural reinstatement tracks spread of attention between object features in working memory. *Journal of Cognitive Neuroscience*, *34*, 1681–1701. https://doi.org/10.1162/jocn_a_01879, PubMed: 35704549
Quax, S. C., Dijkstra, N., van Staveren, M. J., Bosch, S. E., & van Gerven, M. A. (2019). Eye movements explain decodability during perception and cued attention in MEG. *Neuroimage*, *195*, 444–453. https://doi.org/10.1016/j.neuroimage.2019.03.069, PubMed: 30951848
Rac-Lubashevsky, R., & Frank, M. J. (2021). Analogous computations in working memory input, output and motor gating: Electrophysiological and computational modeling evidence. *PLoS Computational Biology*, *17*, e1008971. https://doi.org/10.1371/journal.pcbi.1008971, PubMed: 34097689
Souza, A. S., Rerko, L., & Oberauer, K. (2014). Unloading and reloading working memory: Attending to one item frees capacity. *Journal of Experimental Psychology: Human Perception and Performance*, *40*, 1237–1256. https://doi.org/10.1037/a0036331, PubMed: 24730737
Spitzer, B., & Haegens, S. (2017). Beyond the status quo: A role for beta oscillations in endogenous content (re) activation. *eNeuro*, *4*, ENEURO.0170-17.2017. https://doi.org/10.1523/ENEURO.0170-17.2017, PubMed: 28785729

Sprague, T. C., Ester, E. F., & Serences, J. T. (2014). Reconstructions of information in visual spatial working memory degrade with memory load. *Current Biology*, *24*, 2174–2180. https://doi.org/10.1016/j.cub.2014.07.066, PubMed: 25201683

Sprague, T. C., Ester, E. F., & Serences, J. T. (2016). Restoring latent visual working memory representations in human cortex. *Neuron*, *91*, 694–707. https://doi.org/10.1016/j.neuron.2016.07.006, PubMed: 27497224

Sternberg, S. (1969). The discovery of processing stages: Extensions of Donders' method. In W.G. Koster (Ed.), *Attention and performance II* (pp. 276–315). Amsterdam: North-Holland. https://doi.org/10.1016/0001-6918(69)90055-9

Stokes, M. G. (2015). "Activity-silent" working memory in prefrontal cortex: A dynamic coding perspective. *Trends in Cognitive Sciences*, *19*, 394–405. https://doi.org/10.1016/j.tics.2015.05.004, PubMed: 26051384

Thom, J. L., Nobre, A. C., van Ede, F., & Draschkow, D. (2023). Heading direction tracks internally directed selective attention in visual working memory. *Journal of Cognitive Neuroscience*, *35*, 856–868. https://doi.org/10.1162/jocn_a_01976, PubMed: 36802368

van Driel, J., Olivers, C. N. L., & Fahrenfort, J. J. (2021). High-pass filtering artifacts in multivariate classification of neural time series data. *Journal of Neuroscience Methods*, *352*, 109080. https://doi.org/10.1016/j.jneumeth.2021.109080, PubMed: 33508412

van Ede, F., Checkroud, S. R., & Nobre, A. C. (2019). Human gaze tracks attentional focusing in memorized visual space. *Nature Human Behavior*, *3*, 462–470. https://doi.org/10.1038/s41562-019-0549-y, PubMed: 31089296

van Ede, F., Chekroud, S. R., Stokes, M. G., & Nobre, A. C. (2019). Concurrent visual and motor selection during visual working memory guided action. *Nature Neuroscience*, *22*, 477–483. https://doi.org/10.1038/s41593-018-0335-6, PubMed: 30718904

van Ede, F., & Nobre, A. C. (2023). Turning attention inside out: How working memory serves behavior. *Annual Review of Psychology*, *74*, 137–165. https://doi.org/10.1146/annurev-psych-021422-041757, PubMed: 35961038

Wolff, M. J., Jochim, J., Akyurek, E. G., & Stokes, M. G. (2017). Dynamic hidden states underlying working-memory-guided behavior. *Nature Neuroscience*, *20*, 864–871. https://doi.org/10.1038/nn.4546, PubMed: 28414333

Yu, G., Herman, J. P., Katz, L. N., & Krauzlis, R. J. (2022). Microsaccades as a marker not a cause for attention-related modulation. *eLife*, *11*, e74168. https://doi.org/10.7554/eLife.74168, PubMed: 35289268