Interhemispheric transfer of working memories

Scott L. Brincat*[1], Jacob A. Donoghue*[1] , Meredith K. Mahnke[1], Simon Kornblith[1], Mikael Lundqvist[1,2], and Earl K. Miller[1]

*Indicates equal contribution

[1]The Picower Institute for Learning and Memory and Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, Cambridge, MA 02139 USA

[2]Department of Psychology, Stockholm University, SE-106 91, Stockholm, Sweden

Summary

Visual working memory (WM) storage is largely independent between the left and right visual hemifields/cerebral hemispheres, yet somehow WM feels seamless. We studied how WM is integrated across hemifields by recording neural activity bilaterally from lateral prefrontal cortex. An instructed saccade during the WM delay shifted the remembered location from one hemifield to the other. Before the shift, spike rates and oscillatory power showed clear signatures of memory laterality. After the shift, the lateralization inverted, consistent with transfer of the memory trace from one hemisphere to the other. Transferred traces initially used different neural ensembles from feedforward-induced ones but they converged at the end of the delay. Around the time of transfer, synchrony between the two prefrontal hemispheres peaked in theta and low-gamma frequencies, with a directionality consistent with memory trace transfer. This illustrates how dynamics between the two cortical hemispheres can stitch together WM traces across visual hemifields.

Keywords

working memory, prefrontal cortex, interhemispheric, cognition, neural synchrony

Introduction

Imagine driving on the freeway. A car passes you and holds your attention as you wait to see if it will cut in front of you. Even if you briefly close your eyes or shift your gaze elsewhere, you are able to maintain the car's location in mind and are surprised if it changes unexpectedly. This relies on visual working memory (WM), the ability to maintain images in mind in their absence. Decades of evidence points toward prefrontal cortex (PFC) as a key node in the cortical network underlying WM (D'Esposito and Postle, 2015; Funahashi et al., 1989; Fuster and Alexander, 1971; Miller et al., 1996; Romo et al., 1999; Ungerleider et al., 1998; Voytek and Knight, 2010).

In both humans and monkeys, visual WM seems largely independent between the left and right visual hemifields, which project to the right and left cerebral hemispheres, respectively. WM has

a very limited storage capacity (Luck and Vogel, 1997, 2013). But the capacity within one visual hemifield is largely unaffected by the number of objects in the other hemifield (Buschman et al., 2011; Delvenne, 2005; Umemoto et al., 2010). Correspondingly, the neural correlates of WM storage and WM load (how many items are held in memory) primarily reflect items within the contralateral hemifield (Funahashi et al., 1990; Kastner et al., 2007; Kornblith et al., 2015; Luria et al., 2016; Rainer et al., 1998).

Nevertheless, visual cognition seems seamless across the visual field, even when eye movements switch the remembered location of objects between visual hemifields. In such situations, are memory representations transferred from one cerebral hemisphere to the other or are they “bound” to the hemisphere where they were initially stored? If transferred, what are the underlying mechanisms? We utilized a novel variant of the delayed nonmatch-to-sample task. A midline-crossing saccade during the memory delay switched the hemifield of a remembered item. We found that WM traces are transferred from one prefrontal hemisphere to the other, and that this transfer is facilitated by rhythmic coupling between the cerebral hemispheres.

Results

Monkeys were trained to perform a modified version of a nonmatch-to-sample visual WM task (Fig. 1A). They fixated on a point on the left or right (50% of trials randomly) of a computer screen. An object briefly appeared as a sample in the center of the screen, thus in the right or left visual hemifield, respectively (Fig. 1A, insets). The sample could be one of two different objects, at one of two different locations slightly above or below the center of the screen. The monkeys were required to remember both the object identity and its upper vs lower location over a blank delay, and then compare it to a test object. If it did not match the sample in either identity or upper/lower location, they were trained to saccade directly to it. Otherwise, they were to hold fixation until a second, always non-matching test object appeared.

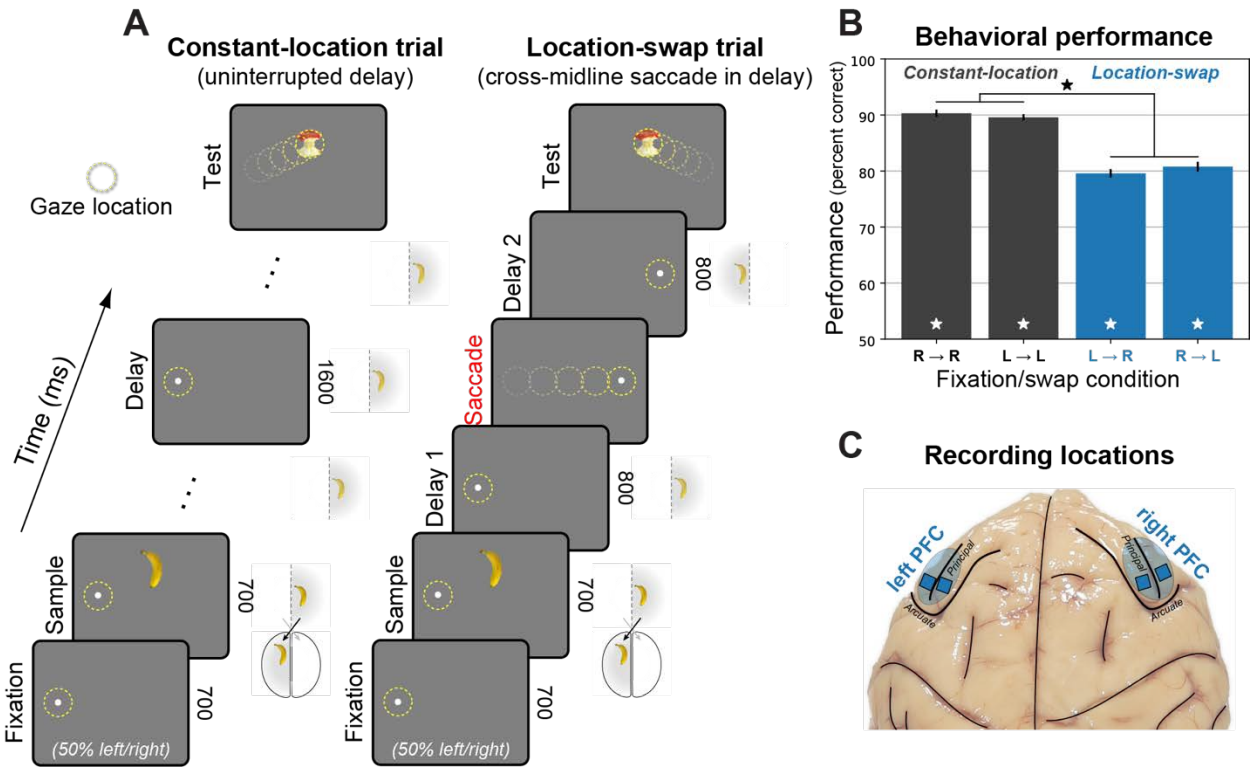

Figure 1. Behavioral and electrophysiological methods. *(A) Hemifield-swap working memory (WM) task. Subjects fixated to the left or right, while a sample object was presented in the center, placing it in the right or left visual hemifield, respectively (insets). Samples could be one of two objects, presented in one of two locations (above or below center). After a delay, a series of two test objects was displayed, and subjects responded to the one that did not match the sample in object identity or upper/lower location (response to first object shown for brevity). In constant-location trials (left), the WM delay was uninterrupted. In location-swap trials (right), subjects were instructed to saccade to the opposite side mid-delay, switching the visual hemifield of the remembered location relative to gaze (insets). (B) Mean performance (± SEM across 56 sessions) for each swap condition and visual hemifield. Monkeys performed the task well (white stars: significant vs chance), with a small but significant decrease (black star) on location-swap trials. (C) Electrophysiological signals were recorded bilaterally from 256 electrodes in lateral prefrontal cortex (PFC).*

A random 50% of trials had an uninterrupted delay (*constant-location* trials because the remembered sample location did not change relative to gaze; Fig. 1A, left). On other trials, halfway through the delay the fixation point jumped across the midline to the opposite side, instructing an immediate saccade, and refixation on it for the remainder of the delay. This shifted the remembered sample's retinotopic location to the opposite visual hemifield (*location-swap* trials; Fig. 1A, right). Performance was good for all conditions (all $p \leq 1\times10^{-4}$, randomized sign test across 56 sessions), albeit somewhat worse on location-swap trials ($p \leq 1\times10^{-4}$, permutation paired *t*-test; Fig. 1B). There were no significant differences in performance when the sample ($p = 0.71$) or test object ($p = 0.10$) appeared in the left vs right hemifield, so all results were pooled across them. We recorded from 256 electrodes in four chronic arrays implanted bilaterally in both hemispheres of lateral prefrontal cortex (Fig. 1C).

Laterality of working-memory-related activity

We first examined constant-location trials to establish a reliable signature of the laterality of the memory trace. Data from each prefrontal hemisphere and sample object hemifield was analyzed separately and then results were pooled based on whether the sample was contralateral or ipsilateral to the recorded hemisphere.

Spiking was higher and more informative for contralateral than ipsilateral objects. Average multi-unit activity (MUA) was significantly higher for contralateral than ipsilateral samples throughout the trial (Fig. 2A, stars; $p < 0.01$, false discovery rate corrected, permutation paired *t*-test). In fact, for ipsilateral samples, MUA was above baseline only during sample presentation and a brief "ramp-up" at the end of the delay (Fig. 2A, dots; $p < 0.01$, corrected, randomized sign test).

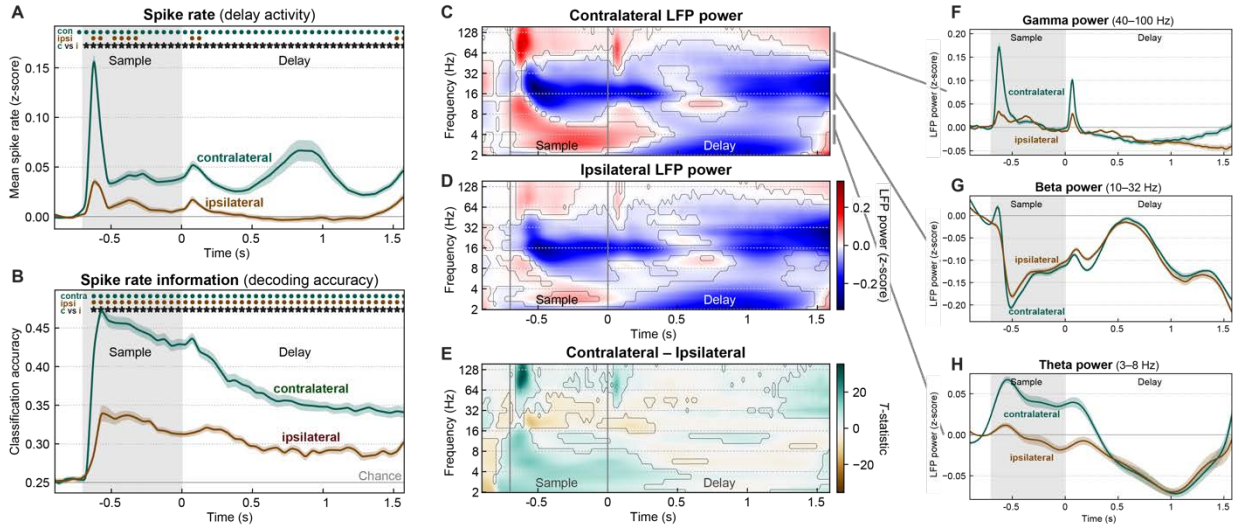

Figure 2. Contralateral bias in prefrontal cortex. *(A) Population mean spike rates (z-scored to baseline, ± SEM across 56 sessions) for sample objects contralateral (green) and ipsilateral (brown) to the recorded prefrontal hemisphere (pooled across left and right). Activity for contralateral samples was greater than baseline (green dots) and greater than activity for ipsilateral samples (stars). (B) Mean (± SEM) information carried in prefrontal spike rates about sample objects (population decoding accuracy) in the contralateral (green) and ipsilateral (brown) visual hemifields. Contralateral information was greater than ipsilateral (stars). (C–D) Mean time-frequency LFP power (z-scored to baseline) for contralateral (C) and ipsilateral (D) samples. Contours indicate significant change from baseline. Gamma (40–100 Hz) and theta (3–8 Hz) power increased from baseline (red), while beta (10–32 Hz) power was suppressed from baseline (blue). (E) Contrast (paired-observation t-statistic map) between contralateral and ipsilateral power. Contours indicate significant difference. (F–H) Summary of LFP power for contralateral (green) and ipsilateral (brown) sample objects, pooled within frequency bands: gamma (F), beta (G), and theta (H). (Note these plots are intended to aid visualization but are not perfect reflections of the full time-frequency responses in C–E due to some time-frequency inseparability of effects.) All modulations from baseline were stronger for contralateral samples, but only gamma showed a difference during the delay period.*

We used a linear discriminant classifier to decode object identity and upper/lower location from the pattern of population activity within each hemisphere at each time point. Cross-validated decoding accuracy was significantly above chance for both contralateral and ipsilateral sample

objects (Fig. 2B, dots; $p < 0.01$, sign test) but it was significantly higher for contralateral (Fig. 2B, stars; $p < 0.01$, paired *t*-test). Thus, prefrontal sensory and memory-related spiking showed a clear contralateral bias, as previously reported (see Discussion).

Local field potential (LFP) power also exhibited a contralateral bias, especially for higher frequencies. Gamma power (~40–100 Hz) was significantly elevated relative to baseline during sample object presentation and a "ramp-up" at the end of the delay for both contralateral (Fig. 2C) and ipsilateral (Fig. 2D) samples (summarized in Fig. 2F; $p < 0.01$, sign test). Gamma power induced by contralateral sample objects was significantly higher than for ipsilateral objects during the sample onset and offset transients and the pre-test "ramp-up" (Fig. 2E; $p < 0.01$, paired *t*-test). Theta power (~3–8 Hz) also showed a contralateral bias during the sample object but not the memory delay (Fig. 2E,H). In contrast, beta power (~10–32 Hz) showed effects in the opposite direction overall—significant *decreases* in power from baseline during the sample and late delay periods for both contralateral and ipsilateral sample objects (Fig. 2G). Like gamma and theta enhancement, beta suppression was significantly stronger for contralateral than ipsilateral samples (Fig. 2E), but, like theta power, only during sample object presentation. These results indicate that prefrontal LFP power, like spiking activity, exhibits a clear contralateral bias.

Transfer of working memories between cerebral hemispheres

We then leveraged the neural signature of WM laterality to examine what happened when the remembered location was shifted by a saccade to the opposite visual hemifield relative to the center of gaze. We propose two alternative hypotheses. On the one hand, a WM trace might be bound to the initial representation induced by visual inputs, and might simply remain in the hemisphere where it was originally encoded (Fig. 3A). This *stable trace model* predicts no change in neural signatures after the mid-delay saccade (Fig. 3B). Alternatively, when the hemifield switches, the neural trace itself might also move from the hemisphere it was originally encoded in to the opposite hemisphere, now contralateral to the remembered location (Fig. 3C). This *shifting trace model* predicts neural signatures of laterality in the location-swap trials (Fig. 3D, light brown and green) will invert after the midline crossing saccade.

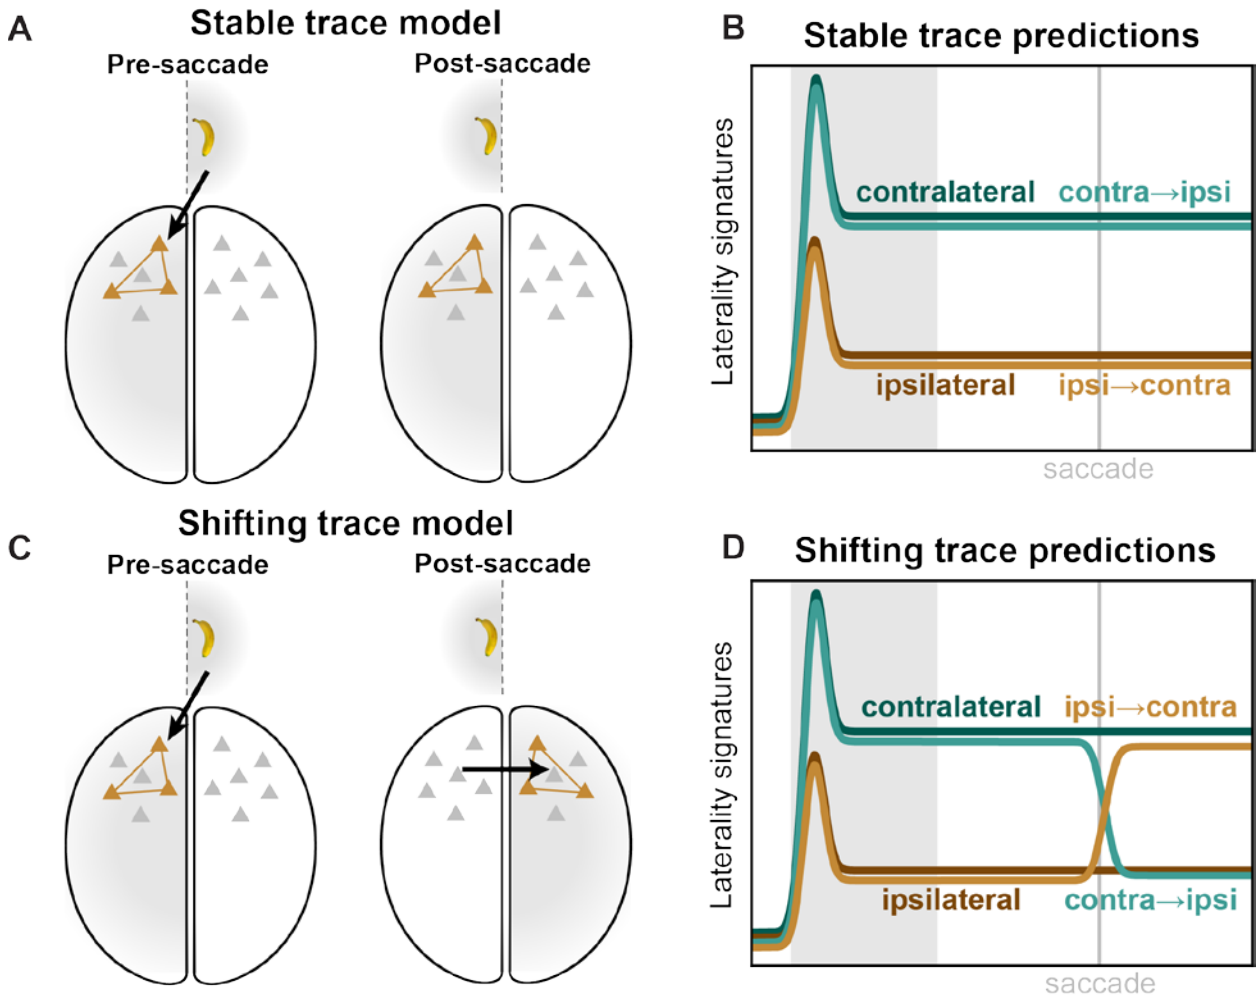

Figure 3. Competing models of location-swap effects. *(A) The stable trace model posits that once a working memory is encoded in a given cortical hemisphere (left), it will remain there (right), despite the remembered location shifting from one hemifield to the other (inset). (B) This model predicts that neural signatures of memory trace laterality will be unaltered by the mid-delay saccade in our task. (C) The shifting trace model assumes that when the hemifield of the remembered location is swapped, the memory trace will be transferred from one cortical hemisphere to the other. (D) This model predicts a post-saccadic inversion of the neural signatures of laterality: shifting the remembered location into the contralateral hemifield (light brown) should come to approximate the constant contralateral location (dark green), while shifting it ipsilateral (light green) should come to look like constant ipsilateral trials (dark brown).*

We found evidence for the shifting trace model, an inversion of neural laterality signatures after the midline-crossing saccade. This was apparent in average MUA rate (Fig. 4A) and in information carried by MUA (Fig. 4B). As in the constant-location trials (Fig. 2A,B), prefrontal MUA starts the delay with a bias toward the contralateral hemifield (Fig. 4A,B, left; 'H' symbols: $p < 0.01$, corrected, sample hemifield main effect in hemifield × shift condition permutation 2-way ANOVA). But after the saccade (Fig. 4A, B, right), both MUA and information increased for remembered locations that shifted from the ipsilateral to the contralateral hemifield (light brown) relative to that for ipsilateral locations on constant-location trials (dark brown; brown stars: $p <$

0.01, paired *t*-test). By contrast, when the saccade shifted the remembered location from the contralateral to ipsilateral hemifield (light green), there was a decrease compared to the contralateral location on constant-location trials (dark green; green stars: $p < 0.01$). For average MUA (Fig. 4A), there was a complete inversion. The spike rates after a saccade that shifted the remembered location into a given hemifield almost exactly matched those for a static memory in the same hemifield. Around the time of the saccade, there was also increased spiking for both location-swap conditions, relative to the constant-location trials ('S' symbols: $p < 0.01$, shift condition main effect). Nevertheless, later in the delay, the predicted inversion effect was dominant ('X' symbols: $p < 0.01$, interaction effect). For information in MUA (Fig. 4B), ipsilateral-shifting trials (light green) were near the value of constant ipsilateral trials (dark brown). Contralateral-shifting trials (light brown) exhibited a bump of increased information after the saccade (brown stars: $p < 0.01$, paired *t*-test) but subsequently declined and never attained the level of constant contralateral trials (dark green). This imperfect transfer of information may explain why behavioral performance was significantly decreased in location-swap trials.

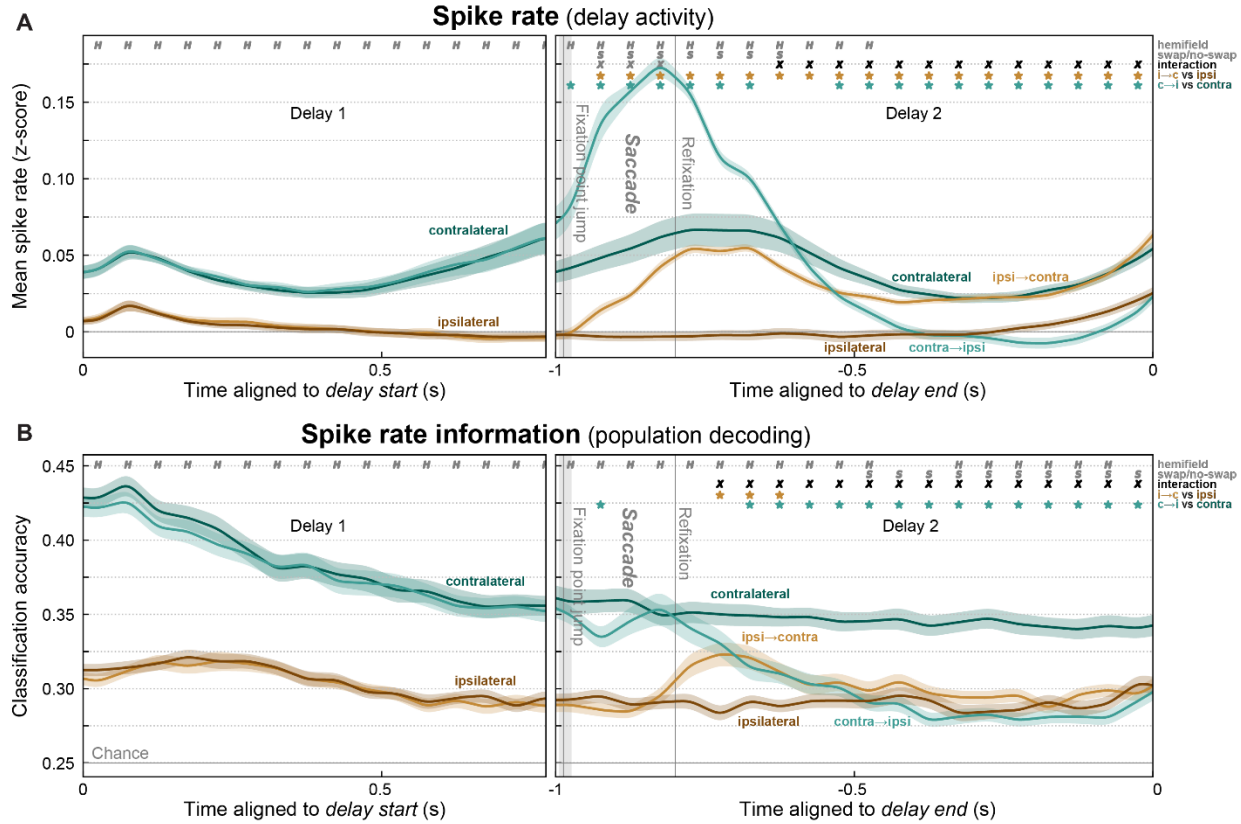

Figure 4. Evidence for interhemispheric transfer of working memory traces in prefrontal spiking activity. *(A) Mean (± SEM) spike rates for all trials where the remembered location was constant in the contralateral (dark green) or ipsilateral (dark brown) hemifield, or where it swapped from ipsilateral to contralateral (light brown) or from contralateral to ipsilateral (light green). Before the mid-delay saccade, there was only a significant effect of the sample hemifield ('H' symbols). Around the saccade, activity was greater overall for location-swap than constant-location trials ('S' symbols). Later, the location-swap trials inverted and approximated activity in the corresponding constant-location trials ('X' symbols: significant hemifield × swap condition*

interaction). Stars indicate significant difference of location-swap conditions from their respective constant-location baseline. (B) Mean (± SEM) information carried in spike rates about the item held in working memory (decoding accuracy). As predicted, post-saccade information decreased in contralateral-to-ipsilateral trials. Information on ipsilateral-to-contralateral trials also significantly increased above baseline, but only transiently, and it never reached the level of constant contralateral trials.

Similar effects were seen in LFP power (Fig. 5). During and just after the saccade, gamma (Fig. 5A–B, D) and theta (Fig. 5A–B,F) power were stronger overall for location-swap than constant-location trials. Later in the delay, however, gamma power inverted and became stronger for remembered locations moving from the ipsilateral to the preferred contralateral hemifield (Fig. 5C,D; $p < 0.01$, hemifield × swap interaction effect) as predicted by the shifting trace model. Theta power showed a similar inversion at the very end of the delay, though this is likely due to temporal smearing of test period effects (Fig. 5C,F).

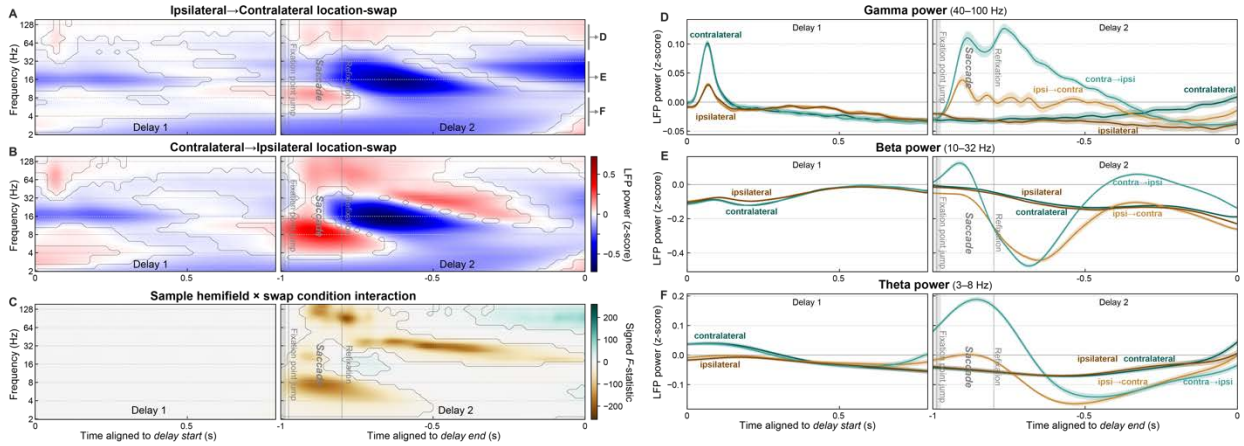

Figure 5. Evidence for interhemispheric transfer of working memory traces in prefrontal LFP power. *(A–B) Mean time-frequency LFP power for trials where the remembered location shifted from the ipsilateral to contralateral (A) or from the contralateral to ipsilateral (B) hemifield. Contours indicate significant difference from pre-sample baseline. (C) Swap inversion effect. F-statistic map for sample hemifield × swap condition interaction, signed to indicate if power was greater when remembered location ends up in contralateral (green) or ipsilateral (brown) hemifield. Contours indicate significant interaction. (D–F) Summary of LFP power for ipsilateral-to-contralateral (light green) and contralateral-to-ipsilateral (light brown) trials, pooled within frequency bands labeled in panel A: gamma (D), beta (E), and theta (F). Around the time of the saccade, LFP power in all bands showed strong effects of the swap condition (saccade vs no saccade). Later in the post-saccade delay, signatures in all bands inverted, as predicted by the shifting trace model.*

Beta power (Fig. 5A–C,E) exhibited complex multiphasic dynamics on location-swap trials (Fig. 5F). It was suppressed initially after the saccade, but later became enhanced, relative to power on constant-location trials (Fig. 5F). On top of these overall dynamics, however, beta power on the location-swap conditions inverted—it became stronger for remembered locations shifting into the ipsilateral hemifield than for those shifting into the contralateral hemifield (Fig. 5C,F). Thus, as for spiking, prefrontal LFP power signatures of working memory laterality also

exhibited the inversion predicted by the shifting trace model (Fig. 3D). These results support the hypothesis that the memory trace is transferred from one cortical hemisphere to the other.

Interhemispheric transfer activates novel neural ensembles

A model consistent with the data thus far is that WMs transferred between hemispheres recruit the same neural ensembles as memory traces activated by feedforward visual inputs into the same cortical hemisphere. Under this *generic ensemble model* (Fig. 6A), when a given object—say, a banana in the upper location—is held in WM within a cortical hemisphere, it uses the same neural ensemble whether it arrived there via feedforward inputs from visual cortex (Fig. 6A, left) or via interhemispheric inputs from the contralateral hemisphere (Fig. 6A, right). We could test this because the saccade on location-swap trials brought the remembered sample location to the same retinotopic coordinates where it appeared on the constant-location trials. An alternative model is motivated by the fact that unique prefrontal ensembles are activated by different combinations of input features and task contexts (Rigotti et al., 2013). Perhaps the same information arriving via different circuits—feedforward (Fig. 6C, left) vs interhemispheric (Fig. 6C, right)—also activates different ensembles. We call this alternative the *novel ensemble* model.

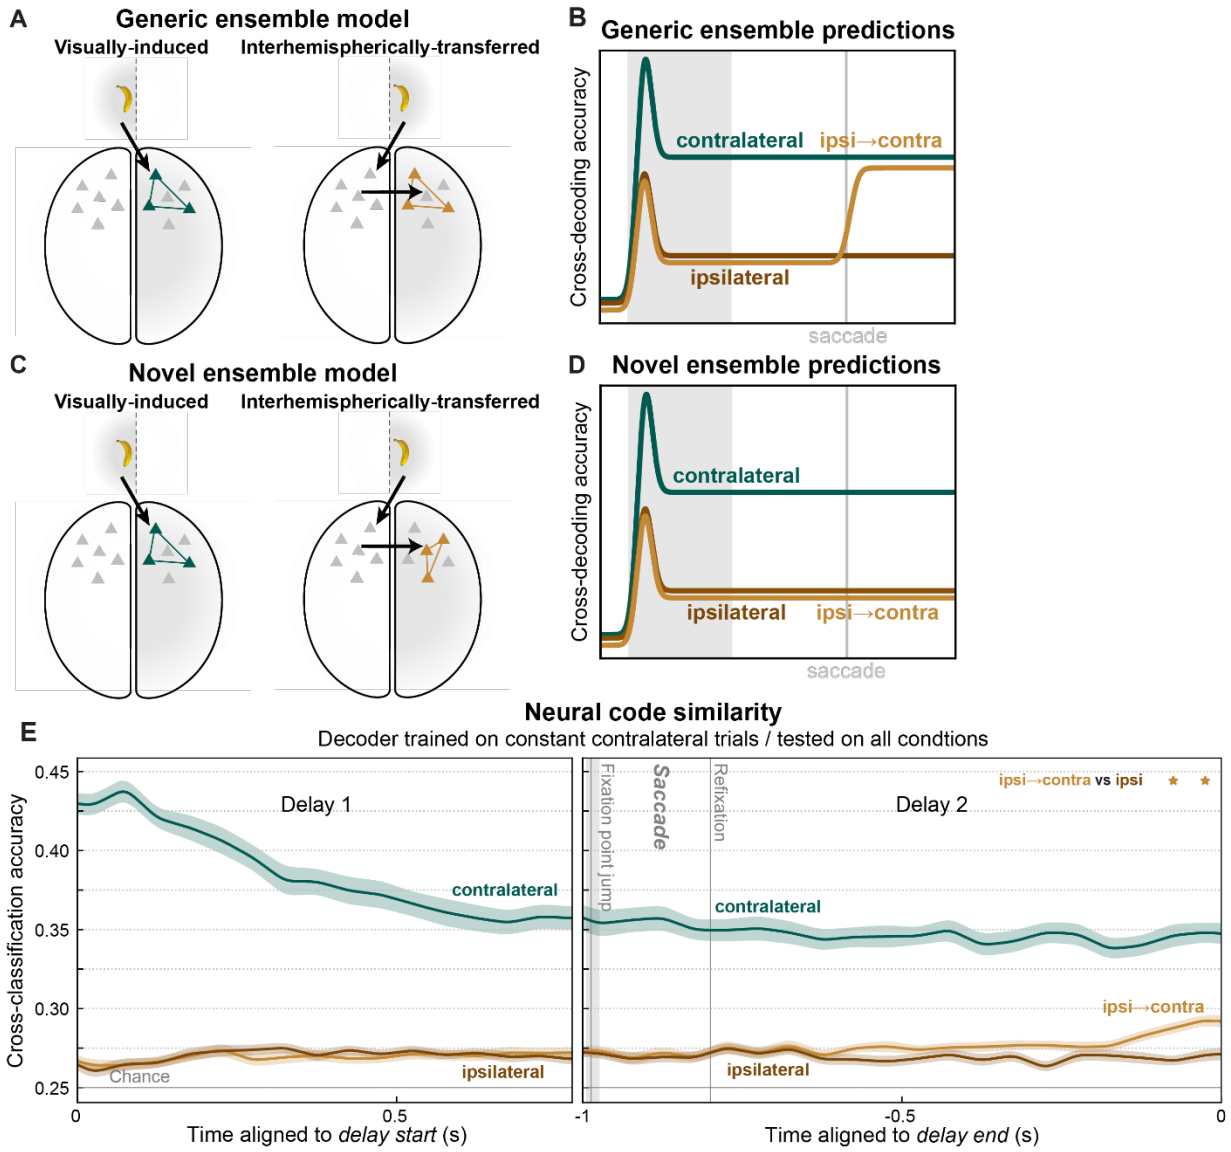

Figure 6. Transferred working memory traces utilized novel ensembles, but converged toward visually-induced ensembles at delay end. *(A) The generic ensemble model assumes a given memory trace will activate the same neural ensemble (colored neurons) whether it arrives in prefrontal cortex via feedforward inputs from visual cortex (left) or via interhemispheric inputs from the opposite cortical hemisphere (right). (B) It predicts a classifier trained to decode working memory traces on constant contralateral trials (dark green) will also be able to decode contralateral-shifting location-swap trials (light brown). (C) The novel ensemble model posits that interhemispheric inputs activate a distinct ensemble (right) from visual inputs (left), even for the same memory trace. (D) It predicts failure of contralateral-trained decoders to generalize to contralateral-shifting trials. (E) For most of the post-saccadic delay, cross-decoding accuracy for ipsilateral-shifting trials (light brown) did not significantly differ from constant ipsilateral trials, as predicted by the novel ensemble model. Near the end of the delay, there was a significant difference (stars), indicating contralateral-shifting trials became more similar to constant contralateral trials, as predicted by the generic ensemble model.*

We used a cross-classification method to adjudicate between the models. We trained a classifier to decode the object identity and its upper/lower location from population spiking at each time point on contralateral constant-location trials. We tested whether these classifiers could predict the same information on ipsilateral-to-contralateral location-swap trials, which brought that same object to the same location as the contralateral constant-location trials. Note that training and testing were both performed on the same cortical hemisphere—separately for each hemisphere, then results were pooled across them—meaning the cross-classification is across *task conditions*, not cortical hemispheres. Thus, we tested whether the same information was reflected in the same neural pattern in a given hemisphere, regardless of how it arrived there. If both conditions activate the same ensembles, as assumed by the generic trace model (Fig. 6A), then this cross-classification (Fig. 6B, light brown) should result in high decoding accuracy, similar to that obtained from both training and testing on constant contralateral trials (Fig. 6B, dark green). If these conditions activate different neural ensembles, as suggested by the novel ensemble model (Fig. 6C), then cross-classification accuracy (Fig. 6D, light brown) should be poor, similar to that obtained from cross-classification testing on constant ipsilateral trials (Fig. 6D, dark brown). We find evidence for both models at different time points during the post-saccade delay.

For reference, we replot (from Fig. 4B) classification accuracy when both training and cross-validated testing were performed on constant contralateral trials (Fig. 6E, dark green). As a control, we computed cross-classification accuracy when the same constant-location contralateral trained classifiers were tested on constant-location ipsilateral trials (Fig. 6E, dark brown). This reflects baseline cross-classification generalization due solely to any bilaterality in the prefrontal neural code. This control analysis resulted in poor accuracy, further evidence for independence between the contralateral and ipsilateral hemifields. For most of the post-saccade delay, cross-classification of the contralateral-shift trials (Fig. 6E, light brown) was also poor and not significantly different from the control ($p \geq 0.01$, paired *t*-test). Here, classifier training and testing were performed at the same time points relative to the end of the delay. Similar results were obtained with training and testing at all possible relative times (Supplement Fig. S*), indicating results are not dependent on the specific timing scheme used.

These results mostly support the novel ensemble model. For much of the post-saccade delay, decoders generalized poorly from the constant contralateral trials to those in which it was transferred from the opposite hemisphere (Fig 6E). The initial bump in information seen when the saccade shifted the remembered location from ipsilateral to contralateral (Fig. 4B) was not detected using the classifier from constant-location trials. Thus, different neural ensembles were activated by the same sensory information in WM, depending on whether it arrived via ipsilateral visual inputs or via the contralateral hemisphere. However, near the end of the delay, the cross-classification of contralateral-shift trials increased relative to the control (Fig. 6E, $p < 0.01$, paired *t*-test). This suggests that, in anticipation of using the WM, interhemispherically-transferred memory traces converge somewhat toward an ensemble representation similar to feedforward-induced traces.

Interhemispheric synchrony during memory transfer

Our results thus far suggest transfer of WM traces between prefrontal hemispheres. If so, we should expect to see evidence of communication between hemispheres around the time of putative transfer. We would further expect signals to flow causally from the hemisphere contralateral to the initial sample hemifield (the "sender") toward the hemisphere contralateral to the post-saccade hemifield (the "receiver"). Evidence suggests phase synchrony between cortical areas helps regulate the flow of information (Fries, 2015). Thus, we measured synchrony between LFPs in the two prefrontal hemispheres using pairwise phase consistency (PPC), an unbiased measure of phase synchrony (Fig. 7).

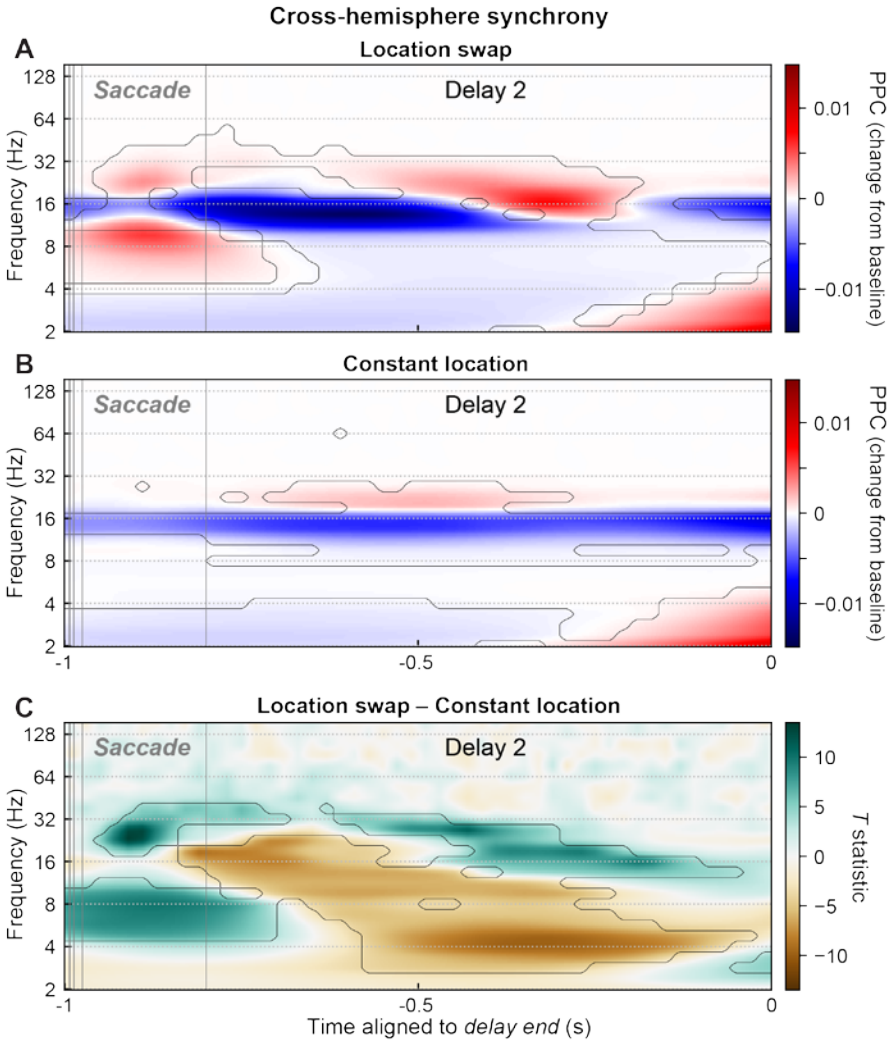

Figure 7. Interhemispheric gamma/theta synchrony may mediate working memory trace transfer. *(A,B) Mean phase synchrony (pairwise phase consistency, PPC) between all pairs of LFPs in the two prefrontal hemispheres, for location-swap (A) and constant-location (B) trials, expressed as the change in PPC from the pre-sample fixation-period baseline. Contours indicate significant differences from baseline. (C) Contrast (paired t-statistic map) between location-swap and constant-location conditions. Contours indicate significant between-condition*

difference. During the time period of putative interhemispheric memory trace transfer (–0.2 to 0 s), there was a significant increase (green) in interhemispheric synchrony in the theta (~4–10 Hz) and beta/low-gamma (~18–40 Hz) bands, and a decrease (brown) in the alpha/low-beta band (~11–17 Hz).

During the saccade on location-swap trials—when WM trace transfer putatively occurs—interhemispheric theta (~4–10 Hz) and high-beta/low-gamma (~18–40 Hz) synchrony both exhibited a transient peak (Fig. 7A; $p < 0.01$, paired *t*-test vs pre-sample baseline). No such peaks were observed at analogous time points on constant-location trials (Fig. 7B). These differences were confirmed by examining the contrast between location-swap and constant-location trials (Fig. 7C; $p < 0.01$, paired *t*-test). In contrast, during this period there was a suppression of interhemispheric synchrony within the alpha/low-beta band (~11–17 Hz), which extended into the post-saccade delay period (Fig. 7C). Similar results, but with higher-frequency gamma effects, were observed for synchrony between LFPs within each hemisphere (Fig. S2). Positive theta and gamma effects predominated in the sender hemisphere (Fig. S2A–C), while suppressive alpha/beta effects predominated in the receiver hemisphere (Fig. S2D–F). These results suggest evidence for signal communication underlying interhemispheric memory trace transfer, and that communication occurs via theta and high-beta/low-gamma—but not alpha/beta—synchrony.

To test whether signals flow from the sender to the receiver hemisphere, we measured spectral Granger causality between LFPs in the two prefrontal hemispheres. This is a measure of how much power at each frequency on one electrode can be explained by power on another, beyond the influence of local dynamics. As predicted, causality was significantly greater in the sender-to-receiver direction (Fig. 8A, green) than in the opposite, receiver-to-sender direction (brown) for all frequencies under approximately 40 Hz (stars; $p < 0.01$, paired *t*-test). This asymmetric directionality was not observed in constant-trace trials between sites contralateral and ipsilateral to the sample location (Fig. 8B). These results suggest signals flow in the predicted direction, from the sender to the receiver hemisphere.

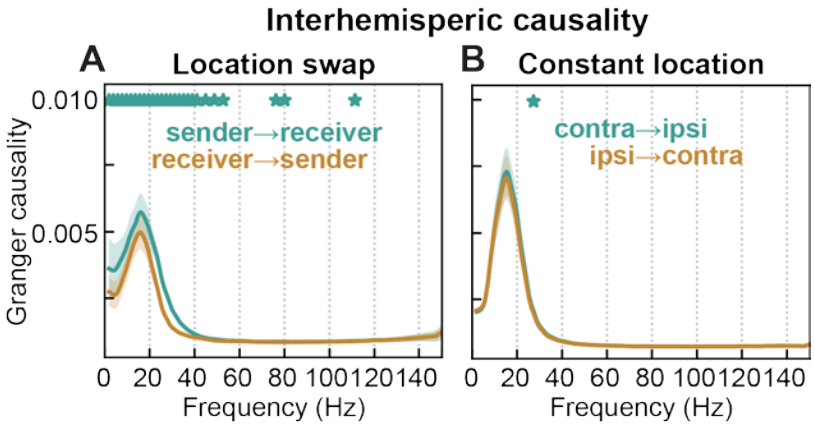

Figure 8. Granger causality flows between prefrontal hemispheres in same direction as putative memory trace transfer. *(A) Mean (± SEM) spectral Granger causality in location-swap trials during time period of putative memory trace transfer between prefrontal hemisphere contralateral to initial sample location ("sender") and hemisphere contralateral to post-saccade*

location ("receiver"). Causality in the sender-to-receiver direction (green) was significantly greater (stars) than in the opposite, receiver-to-sender direction (brown) across all frequencies under ~40 Hz. (B) No asymmetry in interhemispheric causality was observed during analogous time points (relative to delay end) in constant-trace trials between contralateral and ipsilateral hemispheres.

Discussion

Our results suggest that WM traces can be transferred from one prefrontal hemisphere to the other. Previous studies have only provided indirect evidence of this. In WM tasks, when sample and test stimuli appear in opposite visual hemifields, masking during the delay is more effective in the expected test hemifield (Zaksas et al., 2001). Associative cues presented to one visual hemifield can elicit associated representations in the opposite-hemisphere visual cortex (Tomita et al., 1999). Here, we provided direct neurophysiological evidence of WM transfer between hemispheres.

A likely anatomical substrate is direct connections between hemispheres via the corpus callosum (Barbas and Pandya, 1984). They arise mainly from supragranular layer 3 (Schwartz and Goldman-Rakic, 1984), similar to the laminar origin of feedforward projections in cortex (Felleman and Van Essen, 1991). Cortical feedforward processing has been shown to be mediated by gamma and theta oscillations while feedback processing is mediated by alpha/beta oscillations (Bastos et al., 2015; Buschman and Miller, 2007; van Kerkoerle et al., 2014). Similarly, concomitant elevation of gamma and suppression of alpha/beta tracks the encoding and flow of sensory information within prefrontal areas (Lundqvist et al., 2016, 2018). Thus, our results showing increases in interhemispheric theta and high-beta/low-gamma synchrony, and decreases in alpha/beta, during the presumed time of memory trace transfer suggest that interhemispheric processing may rely on similar mechanisms to the feedforward projection of sensory information. The gamma effects we observed (~18–40 Hz) are on the low end of traditional gamma. However, we believe they reflect the same phenomena. Between-hemisphere synchrony (Fig. 7) has a similar pattern of low- and high-frequency enhancement and middle-frequency suppression to that seen in both previous results (Bastos et al., 2015) and our within-hemisphere synchrony (Fig. S2), albeit with the high-frequency effects shifted lower. It may be that communication between cortical hemispheres occurs at somewhat lower frequencies due to the time delays imposed by long callosal axons.

Another way to think about our results is in terms of reference frames. In a retinotopic (gaze-centered) reference frame, locations are described relative to gaze (Supplement Fig. S3A, top). Virtually all studied areas of visual cortex code in a retinotopic reference frame (Fig. S3B, top) (Cohen and Andersen, 2002; Golomb and Kanwisher, 2012). In contrast, a spatiotopic (world-centered) reference frame represents locations in the real-world, independent of gaze (Fig. S3A,B, bottom). In addition to other more established non-retinotopic spatial coding schemes (Chafee et al., 2007; Graziano and Gross, 1998; Olson, 2003), explicit spatiotopic reference frames may exist in higher-level cortex (Dean and Platt, 2006; Duhamel et al., 1997).

In our study, whether the eyes were fixated to the left or right, the sample item was always at the same central spatiotopic (real-world) location. Thus, if prefrontal cortex encoded locations in

a spatiotopic reference frame (Fig. S3B, bottom), its activity should be relatively invariant to the location of the sample object relative to the eyes (Fig. S3C, bottom row). Arguing against this possibility is the differentiation between contralateral and ipsilateral sample locations in PFC activity (Fig. 2). This suggests a retinotopic reference frame (Fig. S3C, top row). On the other hand, if WM—at the cognitive level—also maintained locations in a retinotopic reference frame, remembered locations would be anchored to the location of gaze and would simply shift around with a saccade (Fig. 9A, top). This makes predictions (Fig. 9C, left column) identical to the stable trace model (Fig. 2A,C), i.e., that the memory trace remains in the original hemisphere. Instead, we found that neural signatures of laterality invert after the saccade (Fig. 4–5), ruling out a retinotopic reference frame for WM (Fig. 9C, left column). Our results are consistent with the remembered location being maintained within a spatiotopic reference frame at the cognitive level (Ong et al., 2009) but within a retinotopic reference frame at the neural level (Fig. S3C, upper-right quadrant).

How can we reconcile this? One possibility is that cognition reflects some putative higher-level area with a spatiotopic reference frame that updates the retinotopic representation after each saccade. Arguing against this is the relative paucity of evidence for any explicit spatiotopic representation in the visual cortical hierarchy (Dean and Platt, 2006; Duhamel et al., 1997; Golomb and Kanwisher, 2012). Alternatively, spatiotopic cognition may not rely on an explicit spatiotopic neural representation. It could instead be coded implicitly by a retinotopic representation that is locally updated after a saccade. This would shift a remembered location in the opposite direction of each saccade vector to its new coordinates on the retinotopic map (Pouget and Snyder, 2000). For a midline-crossing saccade, this would entail transfer of the memory trace between the left and right visual hemifields/hemispheres.

In many areas of visual and visuomotor cortex, receptive fields (RFs) show anticipatory shifts to their future, post-saccadic location even before the onset of a saccade (Colby et al., 1995). Other studies suggest that during a saccade, RFs contract toward the location of the saccade target then later expand out to their final post-saccadic location (Chen et al., 2018; Neupane et al., 2016; Zirnsak et al., 2014). This implies a change in the population neural code for location around the time of a saccade because the ensemble responsive to a given location during the RF contraction will be different from those responsive to the same location before and after the saccade. This could explain the lack of cross-decoding between static and shifted memory traces (Fig. 6C). These dynamics have yet to be demonstrated in PFC but their properties in other areas argue against a role. RF contraction effects extend to only ~300 ms after a saccade in area V4 (Neupane et al., 2016) and the frontal eye field (Chen et al., 2018), whereas poor cross-decoding extends to over 500 ms in our results. Contraction effects in V4 occur mainly for neurons with RF in the *same* hemifield as the saccade endpoint (Neupane et al., 2016), whereas in the key ipsilateral-to-contralateral condition in our results, RFs would be in the *opposite* hemifield to the saccade.

An alternative account is motivated by findings of “nonlinear mixed selectivity” in PFC (Rigotti et al., 2013). It suggests PFC is best understood as a random network, in which unique ensembles are activated by different combinations of input features and task contexts (Bouchacourt and Buschman, 2019) . This predicts distinct ensembles are activated depending on the route by

which information arrives, feedforward vs interhemispheric. Regardless of the specific mechanism, our results indicate that just before the memory trace is to be read out for comparison with a test object, its neural code shifts to become more like the ensemble used for static, feedforward-induced memory traces in the same location. This convergence likely facilitates downstream comparison and decision-making processes by allowing similar mechanisms and read-out weights for reading the same information out from WM stores.

Processing in most visual cortical areas, including PFC, is strongly biased toward the contralateral visual hemifield (Funahashi et al., 1990; Hagler and Sereno, 2006; Kastner et al., 2007; Pasternak et al., 2015; Rainer et al., 1998; Voytek and Knight, 2010; Wimmer et al., 2016). Prefrontal spiking activity (Buschman et al., 2011) and gamma power (Kornblith et al., 2015) increase with WM load—the number of items held in memory at one time—but only for items in the contralateral visual hemifield. In contrast, beta power shows increasing suppression for increasing numbers of items in either visual hemifield (Kornblith et al., 2015). Our results confirm these findings. This distinction might reflect the fact that beta oscillations are thought to correlate with broadly-selective inhibitory processes (Engel and Fries, 2010; Jensen and Mazaheri, 2010; Lundqvist et al., 2016). It might also reflect beta having a stronger influence from more bilateral top-down or recurrent signals, while spiking and theta/gamma oscillations are dominated by feedforward signals from strongly lateralized visual cortex (Bastos et al., 2015; van Kerkoerle et al., 2014).

A somewhat surprising result of this lateralization is that WM capacity is largely independent between the two visual hemifields. WM has a very limited capacity for holding multiple items at one time (Luck and Vogel, 1997, 2013). However, in both monkeys (Buschman et al., 2011) and humans (Delvenne, 2005; Umemoto et al., 2010), even when capacity is saturated in one visual hemifield, additional items can be stored in WM if they appear in the opposite hemifield. Similar effects of hemifield independence have been observed with spatial attention (Alvarez et al., 2012) and attentional tracking of moving objects (Alvarez and Cavanagh, 2005). This strong hemifield independence, however, seems inconsistent with the apparently seamless nature of visual WM. Our results provide a possible resolution to this paradox. They suggest that, in such situations, the two prefrontal hemispheres briefly sync up using theta and gamma oscillations in order to physically transfer a WM trace from one cortical hemisphere to its new location on the retinotopic map in the opposite hemisphere.

Acknowledgements

We thank Alexa D'Ambra for assistance, and Jesus Ballesteros, Andre Bastos, Sayak Bhattacharya, Alex Major, Morteza Moazami, Dimitris Pinotsis, and Jefferson Roy for helpful comments. This work was supported by NIMH R37MH087027, ONR MURI N00014-16-1-2832, The MIT Picower Institute Innovation Fund (E.K.M.), and NIGMS T32GM007753 (J.A.D).

Author contributions

J.A.D., S.K., and E.K.M. designed the experiments. J.A.D., M.K.M., and S.K. performed the experiments and recorded the data. S.L.B. curated the data and conceived and performed the analyses. S.L.B, E.K.M., and M.L. wrote the manuscript.

Methods

Behavioral paradigm

Two adult rhesus macaques (*Macaca mulatta*), one male and one female, performed a delayed nonmatch-to-sample working memory (WM) task (Fig. 1A). They began task trials by holding gaze for 700 ms on a fixation point randomly displayed at 4.5° left or right of the center of a computer screen. A sample object was then shown for 700 ms in the center of the screen, thus in the right or left visual hemifield, respectively. Two sample objects were used each session, chosen from a commercial photo library (Hemera Photo-Objects). Sample objects were displayed in one of two positions, 3.4° above or below the screen center. After a 1.6 s delay, a test object was displayed for 400 ms. The monkeys were required to saccade to it if it did not match the remembered sample in either object identity or upper/lower location. If the test was identical to the sample, they withheld response. Then, after a 100 ms blank period, a nonmatching test object was always shown, which required a saccade. Response to the non-match was rewarded with juice, followed by a 3.2 s inter-trial interval.

A random 50% of trials had an uninterrupted 1.6 s WM delay (Fig. 1A, left). In the other 50%, at 800 ms into the delay, the fixation point jumped to the opposite location on the screen (Fig. 1A, right). The monkeys were trained to immediately saccade to it and reacquire fixation. Once fixation was acquired again, the WM delay was continued for another 800 ms, equating the full time of fixated delay period with the constant-location condition.

All stimuli were displayed on an LCD monitor. An infrared-based eye-tracking system (Eyelink 1000 Plus, SR-Research, Ontario, CA) continuously monitored eye position at 1 Hz.

Electrophysiological data collection

The subjects were chronically implanted bilaterally in the lateral PFC with four 8x8 iridium-oxide “Utah” microelectrode arrays (1.0 mm length, 400 µm spacing; Blackrock Microsystems, Salt Lake City, UT), for a total of 256 electrodes (Fig 1C). Arrays were implanted bilaterally in lateral prefrontal cortex (PFC), one array in each ventrolateral and dorsolateral PFC. Electrodes in each hemisphere were grounded and referenced to a separate subdural reference wire. LFPs were amplified, low-pass filtered (250 Hz), and recorded at 30 kHz. Spiking activity was amplified, filtered (250–5,000 Hz), and manually thresholded to extract spike waveforms. All spikes on each electrode were pooled together and analyzed as multi-unit activity (MUA).

Data analysis

General. All correctly performed trials were included in analyses. All analyses of individual MUA and LFP channels were averaged across all electrodes in each hemisphere, and analyses of channel pairs were averaged across all between-hemisphere pairs. Analysis was initially

performed separately for each prefrontal hemisphere and sample object hemifield. Results were then pooled based on whether the sample was contralateral or ipsilateral to a given hemisphere by averaging across appropriate hemisphere/hemifield combinations. This resulted in a set of observations for each experimental session. All plots depict means and standard errors across all 56 sessions in this dataset, and all statistics were performed with sessions treated as observations. All preprocessing and analysis was performed in Python 3.6 or Matlab R2019b (The Mathworks, Inc, Natick, MA).

Preprocessing. Spike rates were computed by binning spike timestamps in non-overlapping 50 ms windows. Spike rates were square-root transformed prior to analysis to convert their Poisson-like distributions to approximately normal. LFPs were re-referenced offline to remove any common-source noise by subtracting off the mean across all electrodes in each array. Evoked potentials were removed by subtracting off the mean signal across trials within each condition (object, upper/lower location, and visual hemifield). Thus, all of our analysis is on the remaining induced component. For most analyses, the resulting signals were convolved with a set of complex Morlet wavelets (wavenumber 6). LFP power was log-transformed to render its distribution approximately normal.

Mean activity analysis. To normalize out any overall differences in activity between neurons and task conditions, we z-scored spike rates to the fixation baseline. Rates were mean-pooled across the 200 ms before sample object onset separately for each neuron and condition, and the mean and standard deviation across all within-condition trials was computed. These were used to z-score rates across all time points and trials for that neuron and condition. The same transformation was used for analysis of LFP power, except it was also computed separately for each frequency, in order to also normalize out the typical 1/*f* distribution of power across frequency.

Population decoding analysis. Spike rates for all multi-units within a given hemisphere were used as independent features in a linear classifier that decoded which of four task conditions—two objects × two upper/lower locations—was present in each trial. Classification was performed independently on spike rate data from each time point (50 ms window). All reported classification accuracies were obtained via 5-fold cross-validation, in which trials were randomly split into five non-overlapping subsets and each classifier was trained on four of these, while its accuracy was evaluated on the final, untrained one. This process was repeated five times with each subset acting as the test set once, and the final results were averaged across the five folds. The same procedure was used for the cross-classification analysis (Fig. 6), except that training and testing trials were selected from different task conditions. For the cross-temporal analysis (Supplementary Fig. S1), cross-classification was also performed using all possible combinations of pairs of time points for training and testing. All decoding analysis was performed with a linear discriminant classifier with optimal covariance shrinkage (Ledoit and Wolf, 2003), using the Python scikit-learn library.

Synchrony analysis. LFP-LFP phase synchrony was computed from the phase of the complex wavelet transform, using the pairwise phase consistency (PPC). PPC is a measure of how consistent across trials the relative phase angles between two signals are, and is an unbiased

estimator of the square of the mean vector resultant length (Kornblith et al., 2015; Vinck et al., 2010).

Causality analysis. Directional causal influences between LFPs in the two prefrontal hemispheres were measured using bivariate nonparametric spectral Granger causality (Dhamala et al., 2008). This quantifies how much power at each frequency in one signal can be explained by power in another, beyond what can be explained by local power. Unlike parametric causality measures, it is estimated directly via factorization of the cross-spectral density matrix, without relying on estimation of a specific autoregressive model. For this analysis only, LFP spectra were computed around the putative trace transfer period (–1 to –0.5 s relative to delay end) via the multitaper method (4 Hz frequency bandwidth, 3 dpss tapers). This analysis was performed using the FieldTrip toolbox for Matlab (Oostenveld et al., 2011).

Postprocessing. To clarify trends in the data, all plotted results were smoothed with 1D or 2D Gaussian kernels for plotting purposes only. The Gaussian standard deviations used were 10 ms for all time axes, and 0.05 octaves for all frequency axes.

Hypothesis testing. All hypothesis tests used non-parametric randomization methods that do not rely on specific assumptions about data distributions (Manly, 2007). Each session was treated as an observation, and randomizations were performed across sessions. All randomization statistics were resampled 10,000 times and evaluated in a two-tailed fashion.

To test whether a mean value differed significantly from baseline, we used a randomized sign test in which a *t*-statistic was computed on both the observed data, and on data where the sign of each baseline-centered observation was randomly flipped. To test whether the means of paired observations were significantly different, we used a permutation paired *t*-test, in which a paired-sample *t*-statistic was computed on the observed data, and on data with the labels of each pair of observations randomly swapped. To test significance of multiple main effects and their interaction, we used a permutation 2-way ANOVA in which *F*-statistics were computed on the observed data, and on data where the multi-factor labels were randomly shuffled as a group across trials. All tests were corrected for multiple comparisons across time points and/or frequencies using a procedure that controls the false discovery rate under arbitrary dependence assumptions (Benjamini and Yekutieli, 2001) using the Python statsmodels module.

Supplemental Information

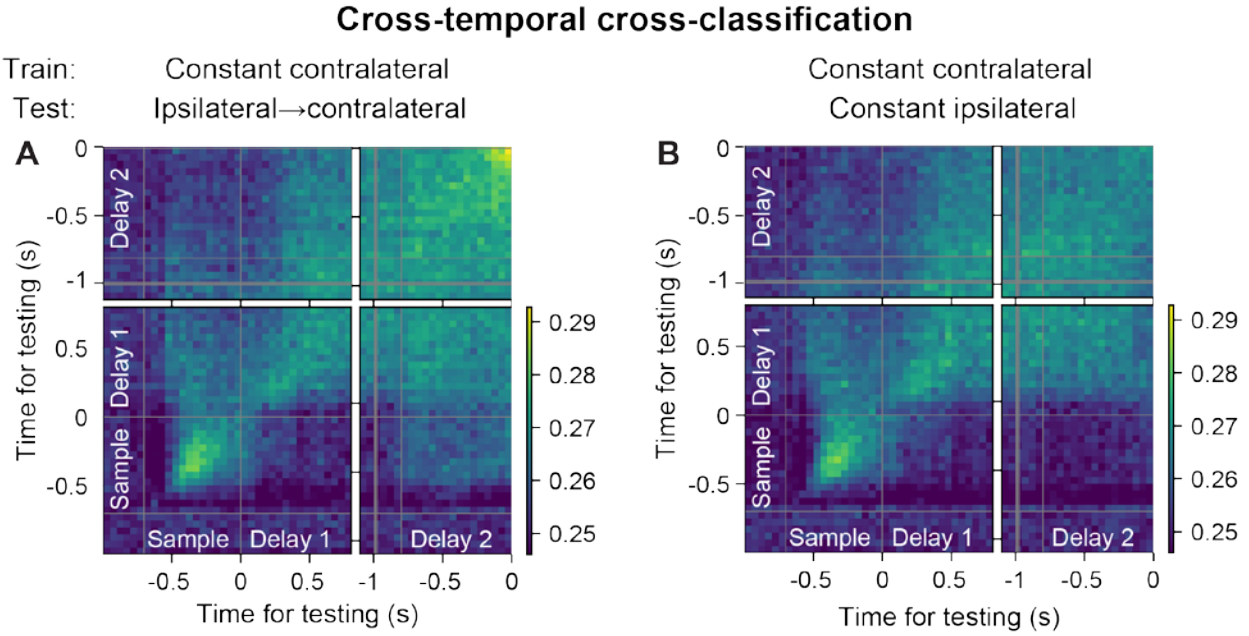

Figure S1. Poor cross-decoding of contralateral shift trials at all relative times. *(A) Mean decoding accuracy for classifiers trained on constant contralateral trials and tested on ipsilateral-to-contralateral shift trials, for all combinations of training time (y-axis) and testing time (x-axis). (B) Cross-temporal decoding accuracy for classifiers trained on constant contralateral trials and tested on constant ipsilateral trials. Main diagonals correspond to training and testing at same time within trials (as plotted in Fig. 6E). This analysis confirms that cross-decoding of ipsilateral-shifting trials (A) is similar to the baseline provided by prediction of constant ipsilateral trials (B) at all training/testing times, except for the "ramp-up" at delay end (upper-right corner of A).*

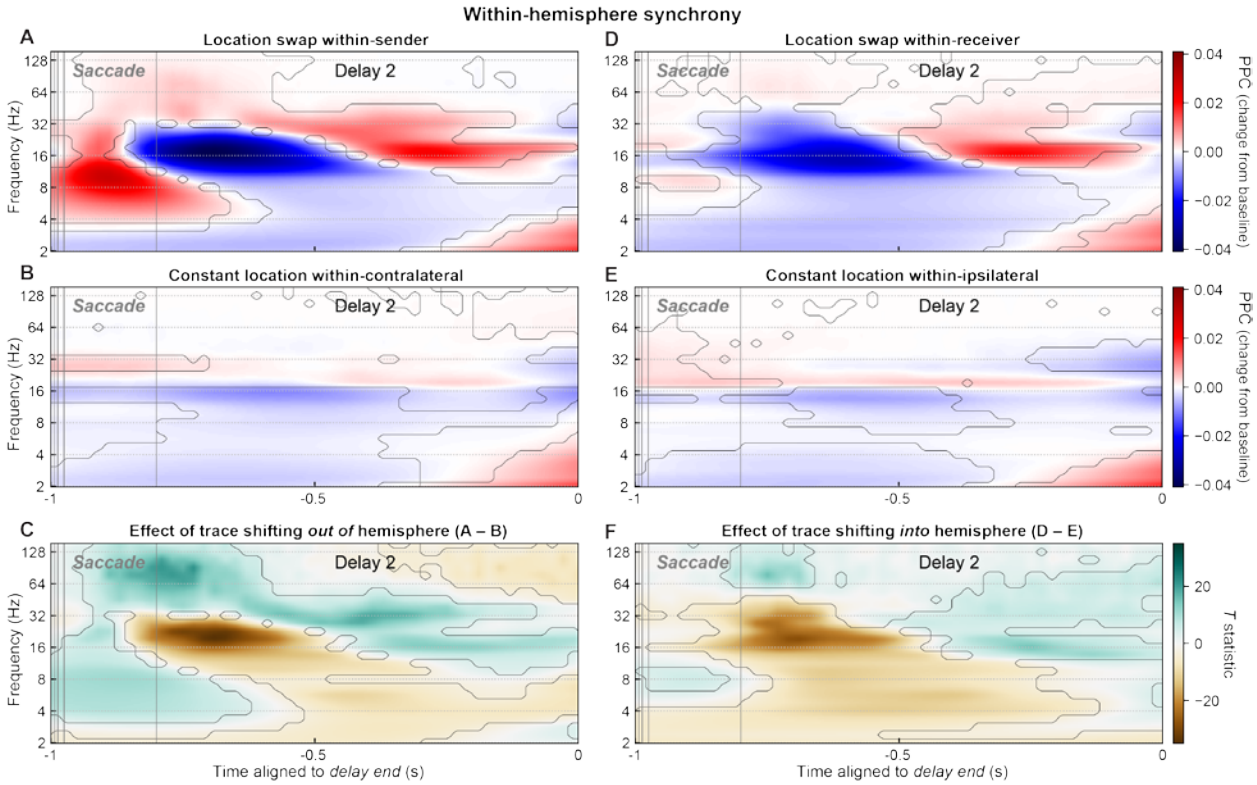

Figure S2. Within-hemisphere synchrony in sender and receiver hemispheres. *(A) Mean phase synchrony (PPC) between all pairs of LFPs within the hemisphere contralateral to the initial sample hemifield (the "sender") on location-swap trials, expressed as the change in PPC from the pre-sample fixation-period baseline. Contours indicate significant differences from baseline. (B) Mean PPC between all LFP pairs within the contralateral hemisphere on constant location trials. (C) Contrast (paired t-statistic map) between (A) and (B), which reveals differences due to a working memory trace shifting out of a hemisphere. Contours indicate significant difference. (D) Mean PPC between all LFP pairs within the hemisphere contralateral to the post-saccade hemifield (the "receiver") on location-swap trials. (E) Mean PPC between all LFP pairs within the ipsilateral hemisphere on constant location trials. (F) Contrast between (D) and (E), which reveals differences due to a memory trace shifting into a hemisphere.*

Around the time of trace transfer (–1 to –0.8 s), synchrony within the sender hemisphere (A,C) shows a similar pattern of increased gamma/theta synchrony and decreased alpha/beta synchrony as between-hemisphere synchrony (Fig. 7A,C in main text), but with the gamma effect at higher frequencies. Receiver-hemisphere synchrony (D,E) shows weaker theta and gamma effects, but beta suppression similar to the sender.

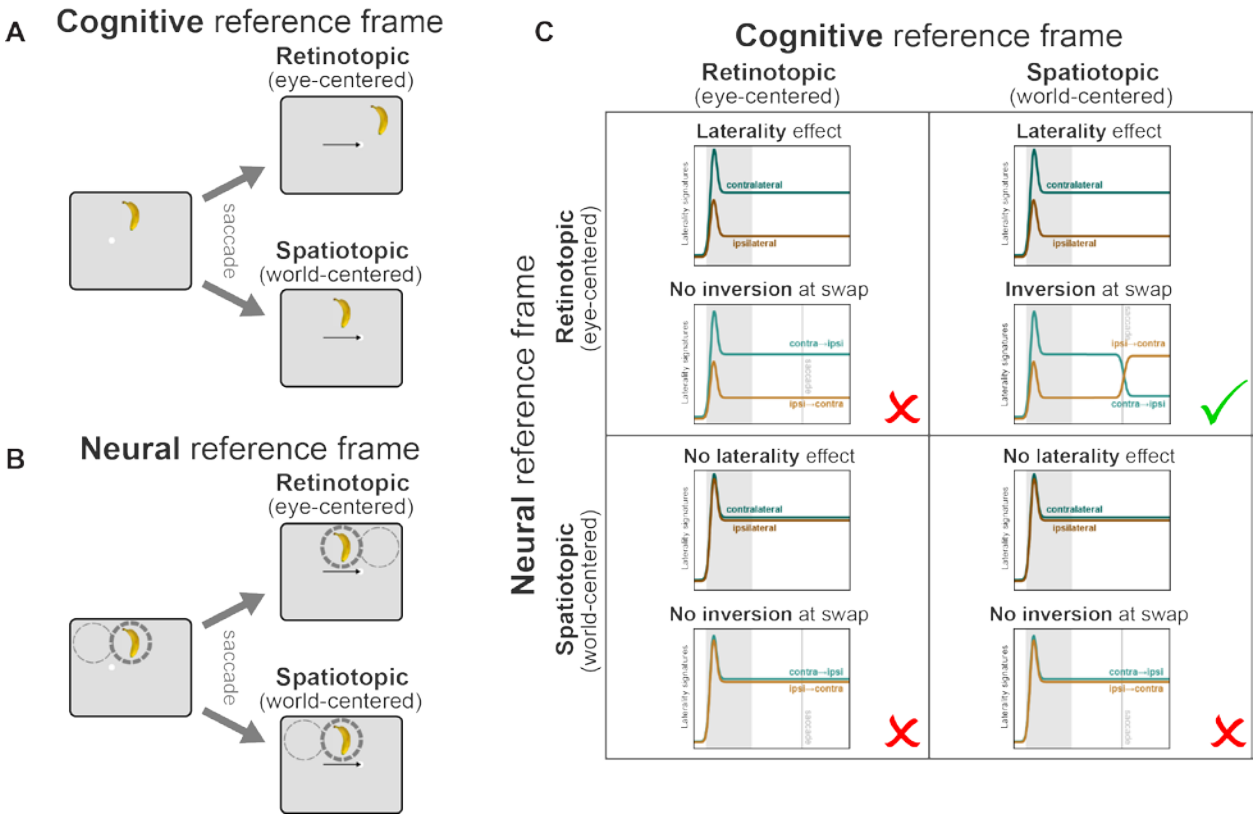

Figure S3. Spatial reference frame interpretation of results. *(A) Illustration of two possible spatial reference frames for remembered locations in cognitive working memory. Left: An object in the right hemifield is encoded into working memory. Right: Effects of a saccade across the midline (black arrow) on the remembered location. Under a retinotopic reference frame (top), the remembered location shifts with the saccade. Under a spatiotopic reference frame (bottom), the remembered location remains anchored to its real-world location. (B) Illustration of two possible reference frames for neural representation of object location. Left: An object in the right hemifield activates the rightmost of two receptive fields (RFs; dashed circles). Right: Under a retinotopic reference frame (top), RFs shift with the eyes, and the object now activates the leftmost RF. Under a spatiotopic reference frame (bottom), RFs are anchored to locations in the real world, and the object continues to activate the rightmost RF. (C) Predictions of all combinations of cognitive (columns) and neural (rows) reference frames for our results. A spatiotopic neural reference frame (bottom row) predicts no change in activity across gaze positions, and thus no laterality effect in our data. This is inconsistent with our results (Fig. 2), ruling out these possibilities. A retinotopic cognitive reference frame (left column) predicts the remembered location shifts with the mid-delay saccade, and thus no inversion of neural laterality signatures. This is inconsistent with our results (Fig. 4–5), ruling out this possibility. Thus, a retinotopic neural reference frame, in conjunction with a spatiotopic cognitive reference frame (upper-right) is the only option consistent with our results.*

References

Alvarez, G.A., and Cavanagh, P. (2005). Independent Resources for Attentional Tracking in the Left and Right Visual Hemifields. Psychol. Sci. *16*, 637–643.

Alvarez, G.A., Gill, J., and Cavanagh, P. (2012). Anatomical constraints on attention: Hemifield independence is a signature of multifocal spatial selection. J. Vis. *12*, 9–9.

Barbas, H., and Pandya, D.N. (1984). Topography of commissural fibers of the prefrontal cortex in the rhesus monkey. Exp. Brain Res. *55*.

Bastos, A.M., Vezoli, J., Bosman, C.A., Schoffelen, J.-M., Oostenveld, R., Dowdall, J.R., De Weerd, P., Kennedy, H., and Fries, P. (2015). Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. Neuron *85*, 390–401.

Benjamini, Y., and Yekutieli, D. (2001). The control of the false discovery rate in multiple testing under dependency. Ann. Stat. 1165–1188.

Bouchacourt, F., and Buschman, T.J. (2019). A Flexible Model of Working Memory. Neuron *103*, 147-160.e8.

Buschman, T.J., and Miller, E.K. (2007). Top-Down Versus Bottom-Up Control of Attention in the Prefrontal and Posterior Parietal Cortices. Science *315*, 1860–1862.

Buschman, T.J., Siegel, M., Roy, J.E., and Miller, E.K. (2011). Neural substrates of cognitive capacity limitations. Proc. Natl. Acad. Sci. *108*, 11252–11255.

Chafee, M.V., Averbeck, B.B., and Crowe, D.A. (2007). Representing Spatial Relationships in Posterior Parietal Cortex: Single Neurons Code Object-Referenced Position. Cereb. Cortex *17*, 2914–2932.

Chen, X., Zirnsak, M., and Moore, T. (2018). Dissonant Representations of Visual Space in Prefrontal Cortex during Eye Movements. Cell Rep. *22*, 2039–2052.

Cohen, Y.E., and Andersen, R.A. (2002). A common reference frame for movement plans in the posterior parietal cortex. Nat. Rev. Neurosci. *3*, 553–562.

Colby, C.L., Duhamel, J.-R., and Goldberg, M.E. (1995). Oculocentric Spatial Representation in Parietal Cortex. Cereb. Cortex *5*, 470–481.

Dean, H.L., and Platt, M.L. (2006). Allocentric Spatial Referencing of Neuronal Activity in Macaque Posterior Cingulate Cortex. J. Neurosci. *26*, 1117–1127.

Delvenne, J.-F. (2005). The capacity of visual short-term memory within and between hemifields. Cognition *96*, B79–B88.

D'Esposito, M., and Postle, B.R. (2015). The Cognitive Neuroscience of Working Memory. Annu. Rev. Psychol. *66*, 115–142.

Dhamala, M., Rangarajan, G., and Ding, M. (2008). Analyzing information flow in brain networks with nonparametric Granger causality. NeuroImage *41*, 354–362.

Duhamel, J.-R., Bremmer, F., Ben Hamed, S., and Graf, W. (1997). Spatial invariance of visual receptive fields in parietal cortex neurons. Nature *389*, 845–848.

Engel, A.K., and Fries, P. (2010). Beta-band oscillations—signalling the status quo? Curr. Opin. Neurobiol. *20*, 156–165.

Felleman, D.J., and Van Essen, D.C. (1991). Distributed hierarchical processing in the primate cerebral cortex. Cereb. Cortex *1*, 1–47.

Fries, P. (2015). Rhythms for Cognition: Communication through Coherence. Neuron *88*, 220–235.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1989). Mnemonic coding of visual space in the monkey's dorsolateral prefrontal cortex. J. Neurophysiol. *61*, 331–349.

Funahashi, S., Bruce, C.J., and Goldman-Rakic, P.S. (1990). Visuospatial coding in primate prefrontal neurons revealed by oculomotor paradigms. J. Neurophysiol. *63*, 814–831.

Fuster, J.M., and Alexander, G.E. (1971). Neuron Activity Related to Short-Term Memory. Science *173*, 652–654.

Golomb, J.D., and Kanwisher, N. (2012). Higher Level Visual Cortex Represents Retinotopic, Not Spatiotopic, Object Location. Cereb. Cortex *22*, 2794–2810.

Graziano, M.S., and Gross, C.G. (1998). Spatial maps for the control of movement. Curr. Opin. Neurobiol. *8*, 195–201.

Hagler, D.J., and Sereno, M.I. (2006). Spatial maps in frontal and prefrontal cortex. NeuroImage *29*, 567–577.

Jensen, O., and Mazaheri, A. (2010). Shaping Functional Architecture by Oscillatory Alpha Activity: Gating by Inhibition. Front. Hum. Neurosci. *4*.

Kastner, S., DeSimone, K., Konen, C.S., Szczepanski, S.M., Weiner, K.S., and Schneider, K.A. (2007). Topographic Maps in Human Frontal Cortex Revealed in Memory-Guided Saccade and Spatial Working-Memory Tasks. J. Neurophysiol. *97*, 3494–3507.

van Kerkoerle, T., Self, M.W., Dagnino, B., Gariel-Mathis, M.-A., Poort, J., van der Togt, C., and Roelfsema, P.R. (2014). Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. Proc. Natl. Acad. Sci. *111*, 14332–14341.

Kornblith, S., Buschman, T.J., and Miller, E.K. (2015). Stimulus Load and Oscillatory Activity in Higher Cortex. Cereb. Cortex bhv182.

Ledoit, O., and Wolf, M. (2003). Improved estimation of the covariance matrix of stock returns with an application to portfolio selection. J. Empir. Finance *10*, 603–621.

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

Luck, S.J., and Vogel, E.K. (2013). Visual working memory capacity: from psychophysics and neurobiology to individual differences. Trends Cogn. Sci. *17*, 391–400.

Lundqvist, M., Rose, J., Herman, P., Brincat, S.L., Buschman, T.J., and Miller, E.K. (2016). Gamma and Beta Bursts Underlie Working Memory. Neuron *90*, 152–164.

Lundqvist, M., Herman, P., Warden, M.R., Brincat, S.L., and Miller, E.K. (2018). Gamma and beta bursts during working memory readout suggest roles in its volitional control. Nat. Commun. *9*.

Luria, R., Balaban, H., Awh, E., and Vogel, E.K. (2016). The contralateral delay activity as a neural measure of visual working memory. Neurosci. Biobehav. Rev. *62*, 100–108.

Manly, B.F.J. (2007). Randomization, bootstrap, and Monte Carlo methods in biology (Boca Raton, FL: Chapman & Hall/ CRC).

Miller, E.K., Erickson, C.A., and Desimone, R. (1996). Neural Mechanisms of Visual Working Memory in Prefrontal Cortex of the Macaque. J. Neurosci. *16*, 5154–5167.

Neupane, S., Guitton, D., and Pack, C.C. (2016). Two distinct types of remapping in primate cortical area V4. Nat. Commun. *7*.

Olson, C.R. (2003). Brain representation of object-centered space in monkeys and humans. Annu. Rev. Neurosci. *26*, 331–354.

Ong, W.S., Hooshvar, N., Zhang, M., and Bisley, J.W. (2009). Psychophysical Evidence for Spatiotopic Processing in Area MT in a Short-Term Memory for Motion Task. J. Neurophysiol. *102*, 2435–2440.

Oostenveld, R., Fries, P., Maris, E., and Schoffelen, J.-M. (2011). FieldTrip: Open Source Software for Advanced Analysis of MEG, EEG, and Invasive Electrophysiological Data. Comput. Intell. Neurosci. *2011*, 1–9.

Pasternak, T., Lui, L.L., and Spinelli, P.M. (2015). Unilateral Prefrontal Lesions Impair Memory-Guided Comparisons of Contralateral Visual Motion. J. Neurosci. *35*, 7095–7105.

Pouget, A., and Snyder, L.H. (2000). Computational approaches to sensorimotor transformations. Nat. Neurosci. *3*, 1192–1198.

Rainer, G., Asaad, W.F., and Miller, E.K. (1998). Memory fields of neurons in the primate prefrontal cortex. Proc. Natl. Acad. Sci. *95*, 15008–15013.

Rigotti, M., Barak, O., Warden, M.R., Wang, X.-J., Daw, N.D., Miller, E.K., and Fusi, S. (2013). The importance of mixed selectivity in complex cognitive tasks. Nature *497*, 585–590.

Romo, R., Brody, C.D., Hernández, A., and Lemus, L. (1999). Neuronal correlates of parametric working memory in the prefrontal cortex. Nature *399*, 470–473.

Schwartz, M.L., and Goldman-Rakic, P.S. (1984). Callosal and intrahemispheric connectivity of the prefrontal association cortex in rhesus monkey: Relation between intraparietal and principal sulcal cortex. J. Comp. Neurol. *226*, 403–420.

Tomita, H., Ohbayashi, M., Nakahara, K., Hasegawa, I., and Miyashita, Y. (1999). Top-down signal from prefrontal cortex in executive control of memory retrieval. Nature *401*, 699–703.

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

Ungerleider, L.G., Courtney, S.M., and Haxby, J.V. (1998). A neural system for human visual working memory. Proc. Natl. Acad. Sci. *95*, 883–890.

Vinck, M., van Wingerden, M., Womelsdorf, T., Fries, P., and Pennartz, C.M.A. (2010). The pairwise phase consistency: A bias-free measure of rhythmic neuronal synchronization. NeuroImage *51*, 112–122.

Voytek, B., and Knight, R.T. (2010). Prefrontal cortex and basal ganglia contributions to visual working memory. Proc. Natl. Acad. Sci. *107*, 18167–18172.

Wimmer, K., Spinelli, P., and Pasternak, T. (2016). Prefrontal Neurons Represent Motion Signals from Across the Visual Field But for Memory-Guided Comparisons Depend on Neurons Providing These Signals. J. Neurosci. *36*, 9351–9364.

Zaksas, D., Bisley, J.W., and Pasternak, T. (2001). Motion Information Is Spatially Localized in a Visual Working-Memory Task. J. Neurophysiol. *86*, 912–921.

Zirnsak, M., Steinmetz, N.A., Noudoost, B., Xu, K.Z., and Moore, T. (2014). Visual space is compressed in prefrontal cortex before eye movements. Nature *507*, 504–507.