Article

Neuron

Gamma and Beta Bursts Underlie Working Memory

Highlights

- Working memory information in neuronal spiking is linked to brief gamma bursts
- The narrow-band gamma bursts increase during encoding, decoding, and with WM load
- Beta bursting reflects a default network state interrupted by gamma
- Support for a model of WM is based on discrete dynamics and not sustained activity

Authors

Mikael Lundqvist, Jonas Rose, Pawel Herman, Scott L. Brincat, Timothy J. Buschman, Earl K. Miller

Correspondence

ekmiller@mit.edu

In Brief

Lundqvist et al. confirm predictions of a working memory model. There is a tight link between gamma oscillations and neural information. Neural events are discrete and thus do not support the modal model of WM based on sustained activity.

Lundqvist et al., 2016, Neuron *90*, 1–13
April 6, 2016
http://dx.doi.org/10.1016/j.neuron.2016.02.028

CellPress

Please cite this article in press as: Lundqvist et al., Gamma and Beta Bursts Underlie Working Memory, Neuron (2016), http://dx.doi.org/10.1016/j.neuron.2016.02.028

Neuron
Article

Gamma and Beta Bursts Underlie Working Memory

Mikael Lundqvist,[1,5] Jonas Rose,[1,2,5] Pawel Herman,[3] Scott L. Brincat,[1] Timothy J. Buschman,[1,4] and Earl K. Miller[1,*]
[1]The Picower Institute for Learning & Memory and Department of Brain & Cognitive Sciences, Massachusetts Institute of Technology (MIT), 43 Vassar Street, Cambridge, MA 02139, USA
[2]Animal Physiology, Institute for Neurobiology, Eberhard Karls University, 72076 Tübingen, Germany
[3]Computational Brain Science Lab, Department of Computer Science & Technology, KTH Royal Institute of Technology, 10044 Stockholm, Sweden
[4]Princeton Neuroscience Institute and Department of Psychology, Princeton University, Princeton, NJ 08544, USA
[5]Co-first author
*Correspondence: ekmiller@mit.edu
http://dx.doi.org/10.1016/j.neuron.2016.02.028

SUMMARY

Working memory is thought to result from sustained neuron spiking. However, computational models suggest complex dynamics with discrete oscillatory bursts. We analyzed local field potential (LFP) and spiking from the prefrontal cortex (PFC) of monkeys performing a working memory task. There were brief bursts of narrow-band gamma oscillations (45–100 Hz), varied in time and frequency, accompanying encoding and re-activation of sensory information. They appeared at a minority of recording sites associated with spiking reflecting the to-be-remembered items. Beta oscillations (20–35 Hz) also occurred in brief, variable bursts but reflected a default state interrupted by encoding and decoding. Only activity of neurons reflecting encoding/decoding correlated with changes in gamma burst rate. Thus, gamma bursts could gate access to, and prevent sensory interference with, working memory. This supports the hypothesis that working memory is manifested by discrete oscillatory dynamics and spiking, not sustained activity.

INTRODUCTION

The ability to keep information available in the absence of sensory input is a key component of working memory (WM) and one of the most studied cognitive functions (Fuster and Alexander, 1971; Goldman-Rakic, 1995; Miller and Cohen, 2001). It is widely assumed to have a neural correlate in sustained neural activity in higher-order cortical areas, such as the prefrontal cortex (PFC) (Fuster and Alexander, 1971; Funahashi et al., 1989; Goldman-Rakic, 1995; Miller et al., 1996; Pasternak and Greenlee, 2005). The mechanism, at first glance, seems straightforward: a sensory event elicits spiking activity that is maintained until that information is needed. This seemingly continuous delay activity may, however, reflect averaging across trials and/or neurons. Closer examination has suggested that the underlying dynamics are more complex (Rainer and Miller, 2002; Shafi et al., 2007; Stokes, 2015). For example, random sampling of neurons indicates that individual neurons bridging a multi-second memory delay is rare. Instead, most neurons show brief bouts of activity with variable onset latency and durations, sprinkled throughout the delay (Cromer et al., 2010; Shafi et al., 2007), suggesting highly dynamic activity (Durstewitz and Seamans, 2006; Stokes et al., 2013).

Continuous, persistent WM information can be simulated by attractor networks, originally serving as models for maintenance of saccade information (Amit and Brunel, 1997; Compte et al., 2000). In these models, information about saccade location is held in a persistent state without interruption. This state corresponds to a dynamic attractor and is supported by recurrent connections that sustain a pattern of activity. If this activity is disrupted, the information it was conveying is lost. By contrast, a related class of attractor models suggests that WM activity is non-stationary. Information is only expressed as spiking during short-lived attractor states. Between the active states, information is held by selective synaptic changes in the recurrent connections and therefore not lost with disrupted activity (Sandberg et al., 2003; Mongillo et al., 2008; Lundqvist et al., 2011, 2012). The limited lifetime of the attractor states has two advantages. First, less spiking is needed to store the information; energy is conserved during the silent states. Second, as information is not lost when activity is disrupted, attractors can hold multiple items in WM with minimal interference between them (or from sensory distractions). In these models, different items are serially encoded and read out, resulting in brief activations of spiking in the coding assemblies.

One of these models (Lundqvist et al., 2011; Figure 1A) implemented the functionality of short-lived attractor states using connectivity and synaptic plasticity constrained by known biology. The model predicts that a burst of gamma oscillations accompanies each attractor state (Figures 1B and 1C) and that the lifetime of such bursts should correspond roughly to an alpha/theta cycle. The gamma oscillations result from fast, local feedback inhibition (Figure 1C), which has two chief consequences. First, firing rates are reduced during attractor retrieval. This state is otherwise characterized by runaway excitation but instead excitation and inhibition are dynamically balanced, leading to the low-rate irregular firing observed in biology (Lundqvist et al., 2010). Second, feedback inhibition normalizes firing rates in a winner-takes-all dynamic, resulting in selective (informative) spiking in only a small subset of neurons (those that are part of the attractor; see Figure 1). This further predicts that there should be a close link

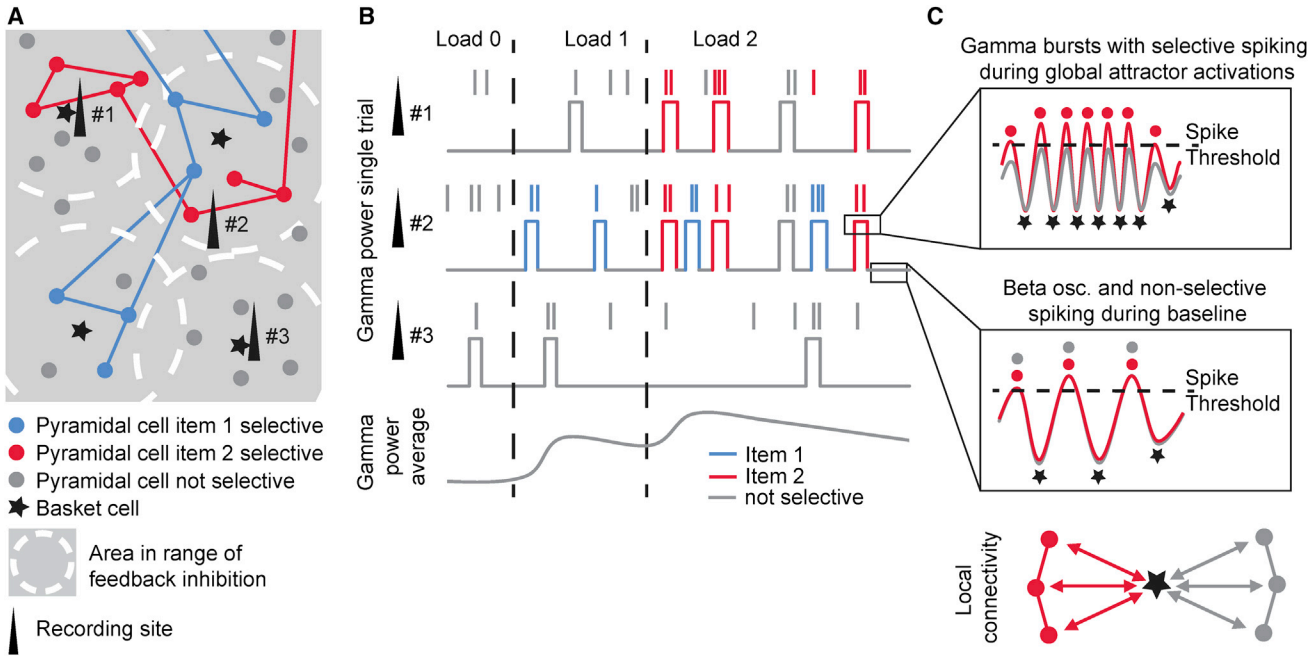

Figure 1. Schematic View of the Model and Model Predictions

(A) Spatial organization of the network. Locally, cells coding for the same stimulus are recurrently connected into local clusters. Several competing clusters share feedback inhibition from nearby inhibitory basket cells. Cell assemblies are formed by recurrent long-range connections connecting several spatially distributed local clusters, each receiving inhibition from a distinct pool of basket cells.

(B) The network displays non-linear attractor dynamics in which various cell assemblies are briefly activated. These activations are initially triggered by stimuli. Following a cell assembly activation, the synapses in the recurrent connections will be potentiated within a certain time window. This will cause the assembly to spontaneously reactivate once it has been triggered by an external stimulus. In this way information about multiple stimuli can be held in WM with attractors, which code for various external stimuli, taking turns in a sequence of reactivations. In the LFP, these activations should be manifested as non-linear transitions into short-lived states with high narrow-band gamma power. These high-power states should become more common in selective sites (but not in non-selective sites, compare sites 1 and 2 with site 3) as WM load increases, leading to enhanced average gamma power with load.

(C) Oscillations (top) are created by local feedback inhibition (bottom) shared by several local clusters of pyramidal cells. During baseline the oscillations are in the beta range and cells from all clusters spike at a similar rate. During attractor activations, there is a slight excitatory bias (from the recurrent connections and assembly specific synaptic potentiation) in one of the assemblies causing this group to consistently reach firing threshold first after each wave of feedback inhibition. As they spike, they activate a new wave of feedback inhibition, shutting down the rest of the cells. Computationally, this creates a winner-takes-all dynamic with spiking only in the (temporarily) most excitable assembly coding for a stimulus. This selectivity in firing implies that the stimulus information conveyed by the corresponding neurons increases. The increased excitation in this state speeds up the oscillations to gamma range.

between information in spiking and gamma power that goes beyond the broad-band increase in gamma power accompanying general increases in spiking activity. The model also predicts that, as more items are stored in the network, they are replayed more and more often leading to a higher density of gamma bursts (Figure 1B; Lundqvist et al., 2011). This could explain observed load-dependent power changes in gamma (Howard et al., 2003; Kornblith et al., 2015; Honkanen et al., 2015), beta (Kornblith et al., 2015; Honkanen et al., 2015), and theta/alpha (Jensen and Tesche, 2002; Palva et al., 2005) in primate cortex.

Non-stationary memory delay activity also has been suggested by observations that PFC activity and gamma oscillations show slow frequency modulation (Jensen and Tesche, 2002; Palva et al., 2005; Watrous et al., 2013; Axmacher et al., 2010). However, the model makes more specific predictions. On a single trial, there should be no prolonged baseline shift in gamma power following stimulus encoding. Gamma power should instead make sharp transitions into the high-power attractor state and repeatedly fall back to pre-stimulus baseline levels throughout the WM delay (thus manifesting what Stokes, 2015, refers to as active-silent states; Figure 1B). As a result, on a trial-by-trial basis, PFC activity is not modulated at slower frequencies in a highly periodic fashion. Instead gamma bursts occur irregularly and the slow periodicity previously observed is instead due to the lifespan of the gamma bursts. The power modulation only appears as periodic when averaging across trials.

We sought to test model predictions in local field potential (LFP) and spike data from the PFC of monkeys performing a multi-item memory task. We did so by performing a unique trial-by-trial analysis of neural activity. This avoided the cross-trial averaging that would obscure the complex temporal dynamics predicted by the model.

RESULTS

We trained two monkeys to retain multiple colored squares over a short memory delay period (Figure 2A). Each trial began with an encoding phase, where two or three squares appeared in a

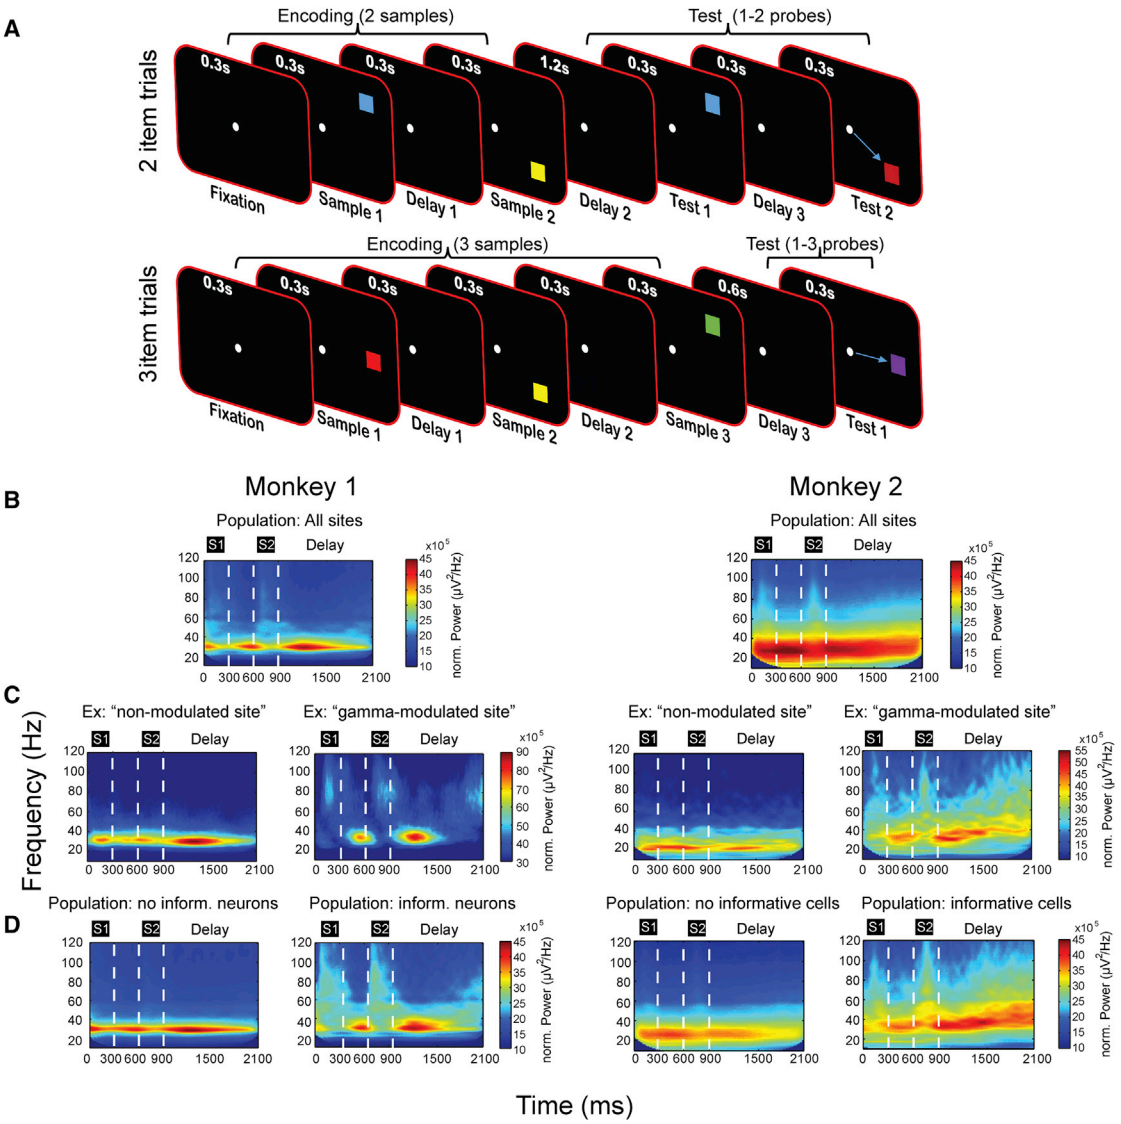

Figure 2. Experimental Setup and LFP Spectral Power

(A) Each trial consisted of three phases: encoding, delay, and test. Following fixation on a white dot in the center of the screen, two (top) or three (bottom) colored squares were sequentially presented to the monkeys. Following a delay period, the stimulus sequence was repeated with the color of one square changed. The monkeys were rewarded for a saccade to the changed square.

(B–D) Spectrograms, not normalized to baseline, of raw LFPs for monkey M1 (left) and M2 (right). The following is displayed for each monkey: (B) average spectrogram from all electrodes and correct trials during encoding (time 0 refers to the onset of the first stimulus) and delay in two-item trials (S1 and S2 refer to samples 1 and 2, respectively); (C) example of spectrograms from single electrodes, recorded the same day, that display non-modulated (left) or gamma-modulated profiles (right); (D) average spectrograms from all electrodes in two-item trials including (right) or excluding (left) neurons that carry information about the presented squares. Power in all spectrograms estimated using multi-taper time-frequency analysis (Experimental Procedures). See also Figures S1, S2, and S3.

sequence, each in a unique location. After a memory delay (1.2 or 0.6 s), there was a test phase in which a new sequence of squares appeared in the same locations as during the encoding phase. However, one of the squares in the test sequence had a different color. The monkey had to respond by making a saccade (Figure S1) to the changed square to receive a reward. Performance was 73% correct on two-item trials and 56% correct on three-item trials, comparable to performance previously reported in monkeys (Buschman et al., 2011). Using multiple acute electrodes in lateral PFC (lPFC) and frontal eye field (FEF), we recorded LFPs and spikes from 321 electrodes with isolatable neurons, which were kept for LFP analysis. Of the isolatable units, we analyzed those with at least four spikes per trial (n = 293).

Prediction 1: Gamma Oscillations Are Tied to Neural Encoding of Information

The tested model predicts specific roles for gamma and beta oscillations. Encoding or decoding a stimulus triggers a gamma

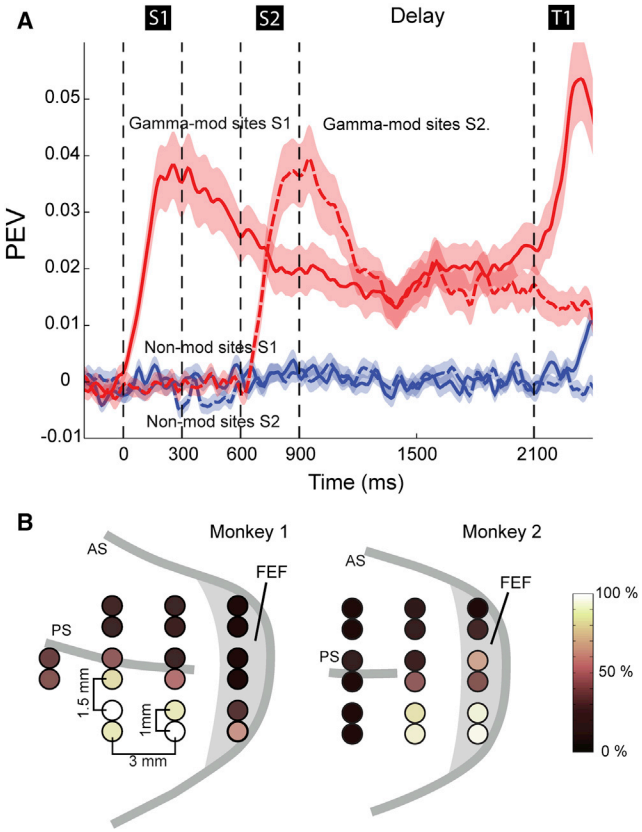

Figure 3. Anatomical Location and Information on Gamma-Modulated Sites

(A) The average information measured using PEV from all cells recorded from gamma-modulated (red) or non-modulated (blue) sites at the time of stimulus presentations (time 0 refers to the onset of the first stimulus). S1 and S2 refer to samples 1 and 2.

(B) The layout of the grid used for inserting the electrodes for monkey M1 (left) and M2 (right). Circles denote recording locations and color code describes the across-session likelihood that an electrode at that particular site displayed increased gamma power during stimulus presentations (gamma-modulated site). See also Figures S4 and S5.

oscillatory attractor state and suppresses a default beta oscillatory state (Figure 1C). This was superficially supported by the average time-frequency evolution for all electrodes (Figure 2B); beta (20–35 Hz) power was suppressed and weak gamma (45–100 Hz) power increases were observed during stimulus presentations. These beta and gamma oscillations were induced (not phase locked to stimuli, Figures S2A and S2B). By contrast, stimulus presentation also evoked lower-frequency (3–10 Hz) activity that was phase locked to stimulus and not present during memory delays (Figures S2A and S2B).

Because of its posited role in gating encoding, gamma should only occur where local neuron populations convey stimulus information. Thus, the model predicts that, at sites where the gamma state is not triggered, spiking is not informative. We first tested this by determining whether gamma was present at some recording sites and not others. Indeed, there was a dramatic difference. The majority of sites (195/321, 61%) showed beta oscillations throughout the trial and no significant increase of gamma in response to stimuli (Figure 2C, referred to as non-modulated sites). However, for the other sites (126/321, 39%; 42.2% monkey M1 and 37.5% monkey M2), gamma power significantly increased following stimulus presentations (Figure 2C, referred to as gamma-modulated sites). A comparison revealed that the largest difference between the gamma-modulated and non-modulated sites was between 45 and 100 Hz (Figure S2C), indicating that gamma power changes indeed reflected oscillations and not the spectral consequences of spiking (investigated in more detail below). In addition, modulated sites showed a stronger beta suppression during stimulus presentations and stronger stimulus-locked 3–10 Hz power (Figure S2C). The latter difference was evident in the evoked potentials (Figure S3), suggesting differences in the processing of bottom-up inputs. Next we determined whether firing in gamma-modulated sites selectively conveyed information about the stimulus.

First, we analyzed the spiking of all recorded neurons. To quantify information they carried about stimuli, we estimated the percentage explained variance (PEV; Olejnik and Algina, 2003; see Experimental Procedures), which measured for each cell the proportion of variance in firing rate that could be explained by stimulus identity (color and location, see Experimental Procedures). About one-quarter of the neurons (24.2%, 71 of 293; see Experimental Procedures for criteria) carried information about the location of the square, its color, or both (Figure S4; average PEV based on color and location combined for informative group was 0.063, PEV based on location only was 0.046). Remarkably, all of these informative neurons (71/71) were recorded at the less common gamma-modulated sites (Fisher's exact test for contingency, $p < 10^{-24}$). The more common non-modulated sites produced spiking activity that did not convey any significant information about stimuli. To further demonstrate this, we recalculated the spectrograms separately for the recording sites that contained any neuron whose spiking conveyed stimulus information (Figure 2D, Population: informative cells) and those that did not (Figure 2D, Population: no informative cells). This yielded virtually identical results to the spectrograms sorted by the presence or absence of stimulus-induced gamma power.

To illustrate this yet further, we plotted the average information from spiking on gamma-modulated versus non-modulated sites (Figure 3A). Neurons from the non-modulated sites had virtually no average stimulus information. They only showed weak information after the test stimulus when the animal made a saccade to the target. By contrast, spiking from gamma-modulated sites showed significant stimulus information starting from sample stimulus onset and continuing throughout the trial (Figure 3A). Thus, induced gamma oscillations and beta-band suppressions at stimulus presentation were strongly co-localized with stimulus information in spiking. This is consistent with the model prediction that stimulus/encoding depends on triggering a gamma attractor state (Figures 1A and 1B) that disrupts the default beta state. If the local neurons do not carry stimulus information, gamma is not triggered and that population stays in the default beta state.

This difference between stimulus information from gamma-modulated versus non-modulated sites was not simply due to

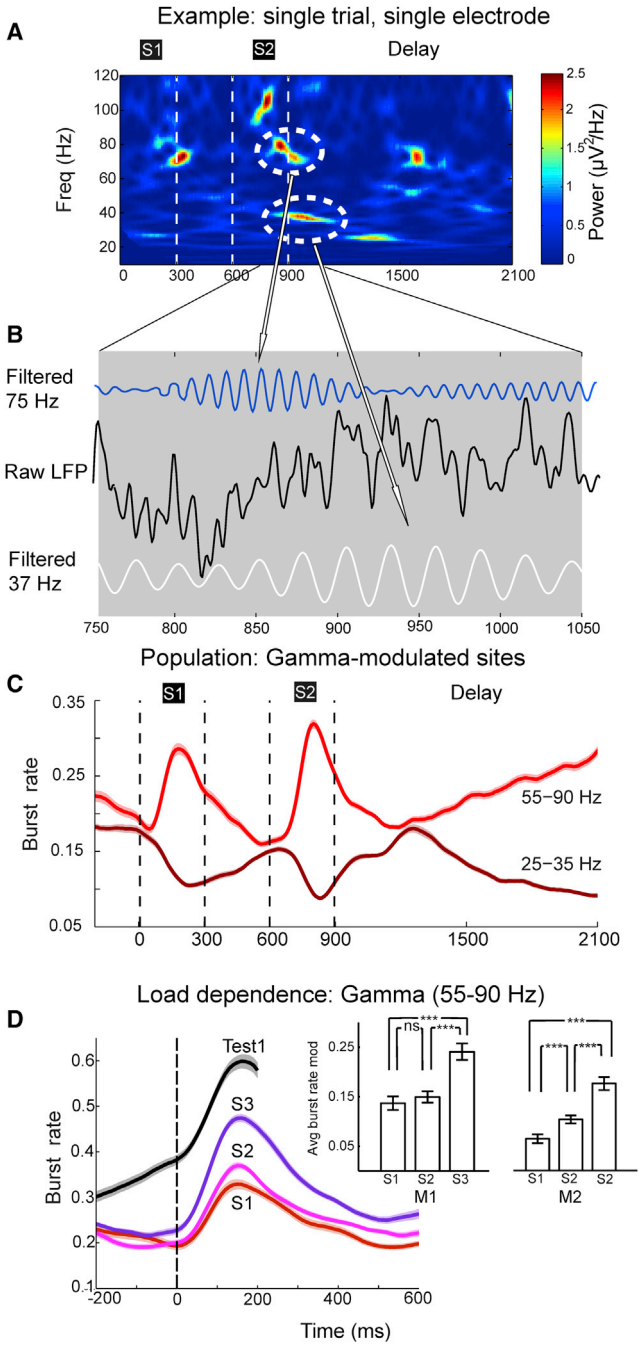

Figure 4. Oscillatory Gamma and Beta Bursts

(A) Example of spectrogram from a single two-item trial.

(B) Zoom-in on the raw LFP (black) around the time of the encircled gamma and beta bursts seen in the top right spectrogram. Blue curve shows the LFP filtered at 75 Hz (center of the encircled gamma burst) and white curve shows LFP filtered at 37 Hz (center of the encircled beta burst.

(C) Estimated burst rate for two-item trials in gamma (light red) and beta (dark red) frequency bands for gamma-modulated sites (both monkeys combined). Burst rate is the time-dependent portion of trials exhibiting a burst at a given time point.

(D) Gamma burst rate in three-item trials around the time of the presentations of each item (S1–S3) and the first test item for the same electrodes as in (C). Inset displays the average size of the gamma burst rate modulation effects in monkey M1 and M2 (significance at level ***$p < 0.01$; error bars denote the SEM). See also Figures S6, S7, and S8.

better quality recordings from the gamma-modulated sites. The average number of isolatable neurons was not higher in the gamma-modulated versus the non-modulated sites (1.41 versus 1.43, $p = 0.81$). The average spiking rate from gamma-modulated sites was only slightly higher than that of non-modulated sites (8.3 versus 6.5 spikes/s; Figure S5). But, to ensure that the differences were not merely due to increased spiking, we gradually removed the neurons with the highest spike rates from gamma-modulated sites until the mean spike rate was the same on the gamma-modulated and non-modulated sites (removing 10% of the most active units was sufficient; Figure S5). This did not change the average information (PEV across all neurons and the entire trial) for the gamma-modulated sites (0.031 versus 0.031, $p = 0.87$), and the information in non-modulated sites remained significantly lower (mean PEV = 0.0006) than in gamma-modulated sites. Similarly, taking 80% of the most active neurons from non-modulated sites resulted in equating average spiking rates between the two types of sites (8.3 spikes/s), but the mean PEV nonetheless remained significantly lower on the non-modulated sites (0.0008; not significantly higher than in the original group, $p = 0.91$). Finally, the gamma-modulated sites were clustered (Figure 3B) in the cortex. As they were sampled with new electrodes each day, the difference between gamma-modulated and non-modulated sites was not due to the electrodes. Instead, there seemed to be an anatomical clustering of WM information where gamma-modulated sites contained a mix of informative and non-informative neurons, while non-modulated sites contained only non-informative neurons (as suggested by the model, compare recording sites 2 and 3 in schematic Figure 1A of the model).

Prediction 2: Gamma and Beta Occur in Brief, Narrow-Band Bursts

The model predicts that, on single trials, beta and gamma power increases should occur in brief, irregular, bursts (Figure 1). Each burst should be narrow band and well defined in frequency, but widely scattered in frequency because it reflects a transient attractor state where the exact frequency is set by the instantaneous level of excitation in the local microcircuit. The average spectrograms in Figure 2 seem to suggest the opposite (broadband, long-lasting signals), but this was an artifact of averaging across trials. Further, the model predicts that, as WM load increases, the number of coding assemblies taking turn being active will increase, leading to more gamma bursts per trial (and consequently higher trial-averaged gamma power; Figure 1). Below we demonstrate that, as predicted, on individual trials both gamma and beta oscillations occurred in brief bursts that were narrow and variable in frequency, like high-power bubbles against low tonic background activity. Examples from a single trial are shown in Figures 4A (spectrogram) and 4B (raw LFP).

To quantify this bubbling, we defined a gamma or beta burst as an increase in power of two SDs above the mean spectral power

in that band, lasting at least three cycles (see Experimental Procedures). The gamma bursting predominantly ranged from 45 to 100 Hz. Each burst on average lasted 67 ms (SD 19 ms) and was narrow in frequency (on average power dropped to 50% relative to the power of the central frequency of the burst within 9.5 Hz, SD 2.9 Hz, see Experimental Procedures). Beta bursting occurred predominately in the 20–35 Hz range, each burst on average lasting 130 ms (SD 37 ms) and also narrow in frequency (mean 5.1, SD 1.7 Hz). To investigate the bursts' relation to trial-averaged power, we calculated burst rate on gamma-modulated sites (Experimental Procedures) by measuring the across-trial density of bursts in each time point (Figure 4C). Both the gamma and beta burst rates were indeed highly correlated with the across-trial power (mean correlation over electrodes: $r = 0.93$, SE = 0.01 for gamma and $r = 0.91$, SE = 0.01 for beta; for all correlations $p < 10^{-8}$; Figure S7). This suggests that the broad-band long-lasting increases in beta and gamma power seen in the average spectrograms were due to gamma and beta bursts of short duration and narrow frequency on a trial-by-trial basis. We confirmed that a bubbling pattern of gamma and beta bursts on single trials gradually turned into smooth continuous broad-band pattern when averaged over an increasing number of trials (Figure S6).

To test the prediction that gamma bursting increases with WM load (Figure 1B), we studied trials with three-item presentations focusing on the time of each sequential stimulus presentation and the first 300 ms of each following delay. Each subsequent stimulus presentation (increase in load; Figure 2A) seemed to result in higher gamma burst rates (Figure 4D). This was tested by quantifying the burst rate over a 150-ms window centered on the peak in average burst rate for each stimulus presentation. Pairwise comparisons between all combinations of stimulus load (i.e., 1 versus 2, 2 versus 3, 1 versus 3) demonstrated significantly higher burst rates for higher loads ($p < 10^{-8}$, Friedman test; pairwise post hoc comparisons with Bonferroni correction at the overall significance $\alpha = 0.05$; see inset for statistics on each monkey separately). Importantly, as also predicted (Figure 1), this was due to an increased number of bursts (burst count during the presentation and delay combined [600 ms] for each load condition progressively increased, $p < 10^{-12}$, testing procedure as above) and not due to load-dependent changes in burst length ($p > 0.31$, Kruskal-Wallis test).

The gamma bursting allowed us to study the relationship between gamma bursts and spiking at a single-trial level. The model predicts that gamma, but not beta, bursts should correspond to states with higher spiking rates (Figure 1C). Indeed, spike rates were significantly higher inside than outside gamma bursts ($p < 10^{-12}$, permutation test on the largest cluster-based statistics; Figure S8A). There was no significant difference between spikes rates inside and outside of beta bursts (Figures S8B and S8C; $p = 0.98$).

The relatively long duration and narrow frequency range of individual bursts quantified above (see also Figure 4A) suggest they are not simply spectral by-products of spiking. Spike-shaped bleeding would have created gamma events with the opposite profile: broad in frequency and even briefer (around 1 ms) (Ray and Maunsell, 2011). In addition, the gamma and beta bursts were visible in raw LFPs (Figure 4B). We further tested whether the gamma oscillations corresponded to rhythmic population activity rather than spike-shaped bleeding by looking at spike-gamma burst interactions as above. Early network models (Brunel and Wang, 2003) with the same pyramidal-interneuron interaction as in the tested model (Figure 1C) suggest that both gamma frequency and spike rates are co-modulated by excitation. We therefore divided gamma into two non-overlapping sub-bands (low: 40–65 Hz versus high: 70–100 Hz). As predicted, the spike rate modulation by gamma bursts was higher in the high gamma than in the low gamma band ($p < 10^{-8}$, Wilcoxon signed-rank test). Thus, higher spike rates are associated with the higher frequency gamma bursts. This would not have been expected if gamma power was due to spike-shaped bleeding, as spike rate then would have had no effect on the spectral component.

Prediction 3: Beta and Gamma Underlie Different Network States

The model further predicts that, because gamma and beta bursting correspond to different states of the network (stimulus coding versus default states, respectively), they should be anti-correlated. The average spectrogram shown in Figure 2D and the plots of burst rates for gamma-modulated sites (Figure 4C; Figure S7) suggest this. Indeed, most of the gamma-modulated sites (112/126) showed a significant negative correlation between gamma and beta power across the trial ($p < 0.05$; $r = -0.41$, SE = 0.04). However, looking into beta (20–25 and 25–35 Hz) and gamma (45–60, 60–75, and 75–90 Hz) power in small 50-ms time bins, we could not find evidence that power in any beta sub-band correlated with power in any gamma sub-band in the same trial and time bin. Thus, there was no apparent instantaneous connection between the two frequencies on a single-trial level.

Prediction 4: Gamma Bursts Are Not Periodic

Several studies (Palva et al., 2005; Axmacher et al., 2010) as well as the model suggest that gamma power should be modulated by slower frequency oscillations. The model explains this low-frequency modulation, not by periodic occurrence of gamma bursts, but by consistency in the duration of the gamma burst per se. To examine this, we analyzed the spectral content of the gamma power envelope. This showed a pronounced peak at 8–10 Hz, which was stronger for gamma-modulated sites and during elevations in gamma burst rate (Figure 5A, compare with Figure 4C). As predicted, this peak was related to the lifetime of gamma bursts rather than their periodic occurrence, as evidenced by the burst auto-correlogram (Figure 5B; Experimental Procedures), which only had a central peak and no side peaks (as would have been expected from a periodic occurrence). In addition, the coefficient of variance (CV_2) for inter-burst intervals was close to 1 (mean CV_2 for delay = 0.997, SD 0.028; mean CV_2 full trial = 0.993, SD 0.018; see Experimental Procedures), indicating a highly variable generation and not periodic appearance. The cross-correlograms also showed a central (albeit much smaller, compare scales) peak (Figure 5C), demonstrating a weak spatial coordination in the timing of bursts across sites.

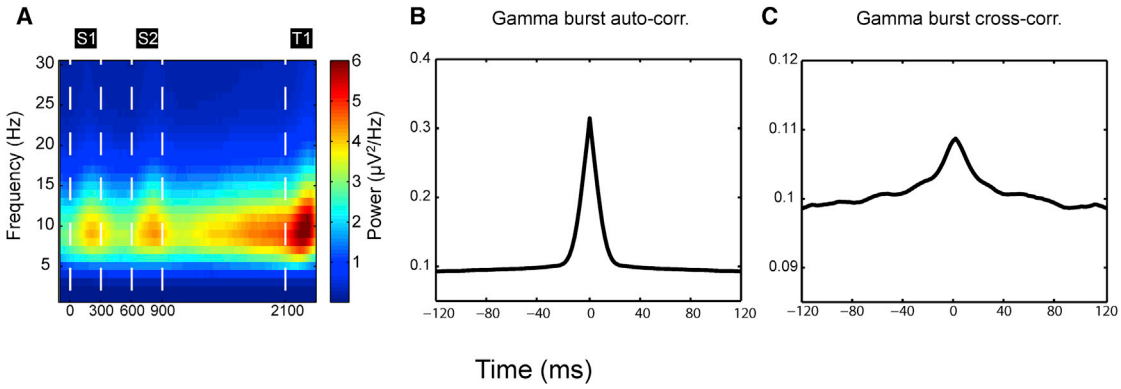

Figure 5. Slow Modulation of Gamma Power

(A) The difference between spectrograms of the envelopes of the gamma-band (45–100 Hz) oscillations in gamma-modulated versus non-modulated sites is shown.

(B) Auto-correlogram of gamma (45–100 Hz) bursts for gamma-modulated sites during the delay period is shown.

(C) Cross-correlogram of gamma bursts calculated between all simultaneously recorded gamma-modulated sites. S1 and S2 refer to samples 1 and 2 and T1 to the first test item.

Prediction 5: Gamma Bursts Are Associated with Stimulus Decoding

The model predicts that gamma bursts not only encode stimuli but also the bursts replay (and beta pauses) when information is read out of WM as spiking. This was supported by our observation that gamma burst rate increased late in the memory delay in two-item (Figure 2C, gamma-modulated sites; Figure 4C) and three-item trials (Figure S7C), accompanied by a drop in beta bursting. This could reflect a replay of stored attractor memories (Lundqvist et al., 2011; Figure 1B) in anticipation of the need to compare them to the test item. This increase in gamma bursting was not related to the anticipation of just any stimulus. Presentations of to-be-encoded sample stimuli at the start of the trial were predictable, but there was no anticipatory increase in gamma bursting (Figure 4D).

To investigate this further, we examined the gamma power evolution on the gamma-modulated sites (with informative spikes; Figure 3A, red lines) versus non-modulated sites (with non-informative spikes). We used spectral band power rather than a burst rate, as the latter measure was inherently normalized across recording sites (Experimental Procedures). The results are illustrated in Figure 6A for both two-item (left panel) and three-item (right panel) trials. In both cases, the gamma power from the gamma-modulated sites (red lines) increased in anticipation of the end of the delay. By contrast, non-modulated sites (without informative spikes, Figure 3A, blue lines) did not show any anticipatory increase. At these sites, the only increase in gamma power occurred at the end of the trial when the animals made their choice, likely due to the saccade. This was also the only period when neurons from the non-modulated sites showed elevated spiking (Figure S5) and carried stimulus information (Figure 3A).

In contrast to gamma, beta-band power in two-item trials dropped at the end of the delay (prior to T1 relative to early delay in Figure 6B), regardless of the site type (gamma-modulated or non-modulated site, $p < 10^{-18}$ for both, Wilcoxon signed-rank test for the average band power between the first and second halves of the delay). Toward the end of delay and during test stimulus, average beta power was not different between gamma-modulated and non-modulated sites ($p = 0.51$, Wilcoxon signed-rank test). This was in contrast to sample presentation when beta was globally suppressed, but more strongly in gamma-modulated sites (Figure 6B; Figure S2; especially relative to the pre-stimulus levels).

To further investigate the role of gamma in encoding and decoding, we examined neurons that carried information about the sample stimuli. There were two subgroups, selected by when in the trial their peak firing rates occurred. For one group of neurons, the peak in firing rates took place during stimulus presentation (the encoding group; Figure 7A, left panel). The other group of neurons showed a peak of spiking during the memory delay (the maintenance group; Figure 7A, right panel). Importantly, the encoding neurons were suppressed at the onset of delay but ramped up activity during the second half of the delay, a property for which they were not selected. Trial-averaged activity of the encoding group, therefore, exhibited a similar trend to the gamma burst rate (Figure 7A, left, red line, gamma-modulated sites). Activity of the maintenance group, by contrast, showed a different temporal profile than the average gamma burst rate (Figure 7A, right, red line). Thus, one group of neurons fired during encoding and decoding epochs, closely following the gamma burst rate, while the other fired when the information needed to be maintained but not accessed. As expected, the time course of neural information in these two groups (PEV, see Experimental Procedures) was strongly correlated with the spiking rate (Figure 7B). Thus, late in the delay, gamma bursting increased ($p < 10^{-14}$) as did spike rate ($p < 10^{-3}$) and information (PEV, $p < 10^{-3}$) in the encoding/decoding group (first 300 ms of the delay versus last 600 ms, Wilcoxon signed-rank test).

It was possible that the increases in gamma bursting and information in the encoding neurons could reflect anticipation of encoding of the test stimulus rather than decoding information from WM per se. To test this, we needed a task in which the monkeys could anticipate the end of a memory delay (and thus the

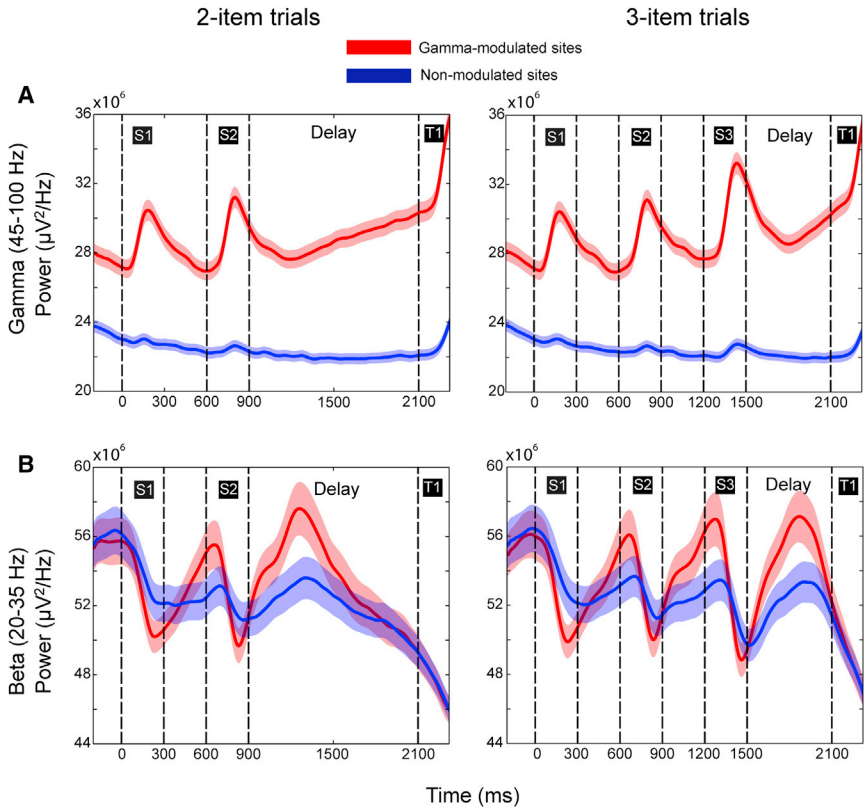

Figure 6. Power in Gamma-Modulated and Non-modulated Sites

(A) Gamma band power in two- (left) and three- (right) item trials averaged for gamma-modulated (red) and non-modulated (blue) sites.

(B) Same as (A), but for beta-band power. Shaded regions represent SEs.

need to decode) without anticipation of a (test) stimulus. To this end, we trained two new monkeys on a spatial delayed-response task (Funahashi et al., 1989; Pesaran et al., 2002; Goldman-Rakic, 1995). The monkeys had to fixate at a dot in the middle of the screen while a target location was flashed (Figure 8A). After a delay period, signaled by the removal of the fixation dot, the animals were to saccade to the location of the previously flashed target. While this task also involved short-term memory maintenance, it was markedly different from the main task. It was less complex and more related to the maintenance of a motor plan rather than the maintenance of two or more stimuli. Critically, at the end of delay, there was no new stimulus as in our original WM task, but instead the cue to use the WM information was the disappearance of the fixation dot. These differences allowed us to test the generality of the anticipatory effects found in the main task.

To more directly verify the hypothesis that gamma bursts are induced as the animal anticipates the need to decode information from WM, we used two different lengths of memory delay. In six out of seven trials the delay was 750 ms while in one out of seven trials the delay was 1,500 ms. This setup was inspired by the observation that gamma bursts seemed to occur earlier in two-item trials, which contained a longer delay interval and were also less common than the three-item trials. We predicted increased gamma and reduced beta burst rates at the end of the delays and that these effects should be particularly strong in the second half of the unexpectedly long delays. This is indeed what we found (Figure 8B; results pooled across all electrodes in lPFC and FEF of both monkeys). The increase in gamma burst rate toward the end of the 750-ms delay was maintained throughout the entire second half of the 1,500-ms delay (1,100–1,850 ms in Figure 8B), and it was accompanied by a decrease in beta burst rate.

DISCUSSION

Model predictions were borne out by trial-by-trial analysis of PFC activity during a multi-item WM task. There was an interplay between spiking activity and induced oscillations in two main frequency bands, beta (20–35 Hz) and gamma (45–100 Hz). Rather than an overall change in the activity state, there were discrete bursts of beta and gamma LFP oscillations, each brief in duration and narrow in frequency range, like small bubbles of oscillatory events. It was only by averaging across trials and recording sites that we observed the broad-band, smoothly varying oscillatory power that has been reported previously (Sederberg et al., 2003; Howard et al., 2003; Roux et al., 2012; Roux and Uhlhaas, 2014; Honkanen et al., 2015; Kornblith et al., 2015). The distinction that the gamma oscillations were narrow band on a single-trial level is of interest, as narrow-band gamma in visual cortex is related to sensory processing (Gray and Singer, 1989; Ray and Maunsell, 2011; Lachaux et al., 2005; Fisch et al., 2009) and has been shown to occur in bursts (Tallon-Baudry and Bertrand, 1999; Kucewicz et al., 2014).

Previous studies reported various different oscillatory correlates of WM (Tallon-Baudry and Bertrand, 1999; Howard et al., 2003; Palva et al., 2005; Jensen and Tesche, 2002; Kornblith et al., 2015). Here we examined spiking and LFPs in tandem. This revealed a close link between informative spiking and gamma bursts. Spiking that conveyed stimulus information was only found at the minority of recording sites that showed increased gamma bursts to stimuli. Given their length (average 67 ms) and narrow within-burst frequency range, the gamma bursts were not the result of spike-shaped bleeding into the high-frequency spectral content, but instead they reflected true oscillatory events (Nir et al., 2007). The gamma bursts and related spiking were induced during encoding and then spontaneously re-appeared throughout the trial, particularly in anticipation of decoding of WM content.

These results are consistent with the tested multi-item WM model (Figure 1). In this model, a brief gamma burst accompanies the coordinated activation of an ensemble coding a

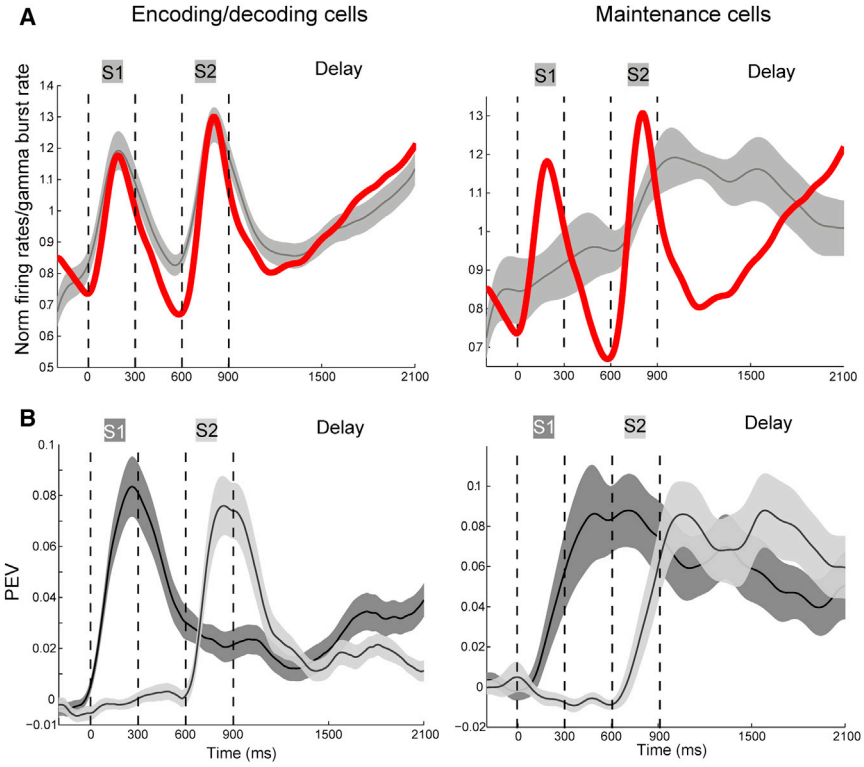

Figure 7. Information and Firing Rates in Two-Item Trials

(A) The average normalized firing rates of all informative early (encoding/decoding, left, n = 50) and late (maintenance, right, n = 21) responding cells. The red curve is the normalized gamma-band (55–90 Hz) burst rate as a reference.

(B) The average PEV based on first (dark) and second (light) stimuli for informative early (encoding/decoding, left) and late (maintenance, right) responding cells. Shaded areas represent SEs.

specific memory item. In particular, the gamma bursts are manifestations of attractor states that correspond to different memory items. Due to the brief lifetime and synaptic foundation of these attractor states, information about distinct items is not retrieved simultaneously, preventing unwanted interference when more than one object is stored (Figure 1). Our findings thus provide support for the discrete coding and periodic replay hypothesis (Sandberg et al., 2003; Lundqvist et al., 2011; Fuentemilla et al., 2010). In this view, WM information is only present in brief bursts of spiking and maintained in synaptic changes between such events (Sandberg et al., 2003; Mongillo et al., 2008; Lundqvist et al., 2011, 2012; Stokes, 2015).

The tested model further predicted that increasing WM load should result in a greater rate, but not length, of gamma bursts (Lundqvist et al., 2011) as more items take turn being active. This prediction is opposite to the main conjecture of a related model (Lisman and Idiart, 1995), where the duration of bursts, but not their number, should increase with load. Both sets of predictions could potentially explain reported increases in mid-range gamma-band power with stimulus load (Howard et al., 2003; Roux et al., 2012; Honkanen et al., 2015; Kornblith et al., 2015). These earlier reports displayed broad-band and relatively long-lasting effects, but our observations suggest that they were likely to be the result of trial averaging. Our trial-by-trial analysis revealed that increased memory load was linked to the increased rate, but not length, of brief narrow-band gamma bursts. The consistent lifetime of bursts also accounted for the slow modulation of average gamma-band power at 8–10 Hz. It was not a continuous periodic modulation, but rather the result of irregular occurrences of bursts of stereotyped length, as predicted by the model. There was no evidence for low-frequency oscillations during the memory delay. This does not preclude a role for slower oscillations, especially in other brain areas (Jensen and Tesche, 2002; Palva et al., 2005; Axmacher et al., 2010; Watrous et al., 2013).

We found that the heterogeneous population of spiking neurons carrying stimulus information could be reduced to two principal populations with different relationships to oscillatory dynamics. One population was mainly active at stimulus encoding and decoding. Its spiking activity profile closely followed the gamma burst rate recorded at the gamma-modulated sites with informative spiking. A second population of informative spiking neurons were active mainly during memory delays when instead the average beta burst rate was higher. Thus, different modes of WM could be reflected by a shift in the balance of beta and gamma burst rates. Activation or suppression of gamma bursting (which was anti-correlated with beta bursting) could activate or suppress spiking of encoding/decoding neurons and, thus, gate access to WM or protect WM from interference from new sensory inputs.

These gamma and beta dynamics may play similar roles in other behaviors and other cortical areas. We observed the same relationship between gamma and beta (and encoding/decoding versus maintenance) in a simpler WM task that involved planning a delayed saccade (Funahashi et al., 1989) and, hence, a motor plan. This is consistent with observations of increased gamma power in lateral intraparietal cortex (LIP) in monkeys during the encoding and decoding epochs of a delayed saccade task (Pesaran et al., 2002). Fast gamma-like oscillations also appear just prior to the onset of movements in M1 (Donoghue et al., 1998). Gamma oscillations often have been associated with stimulus processing (Fries, 2009) and visual attention in sensory cortex (Gregoriou et al., 2009). Further, general beta power suppression, together with specific patterns of increased gamma power (Lachaux et al., 2005; Fisch et al., 2009) or gamma bursting (Tallon-Baudry and Bertrand, 1999; Kucewicz et al., 2014), has been tied to object and category recognition. Indeed, the WM model tested here is an adaptation of a perceptual model (Lundqvist et al., 2006). In that model, partial and noisy inputs activate an internal time-limited attractor state representing

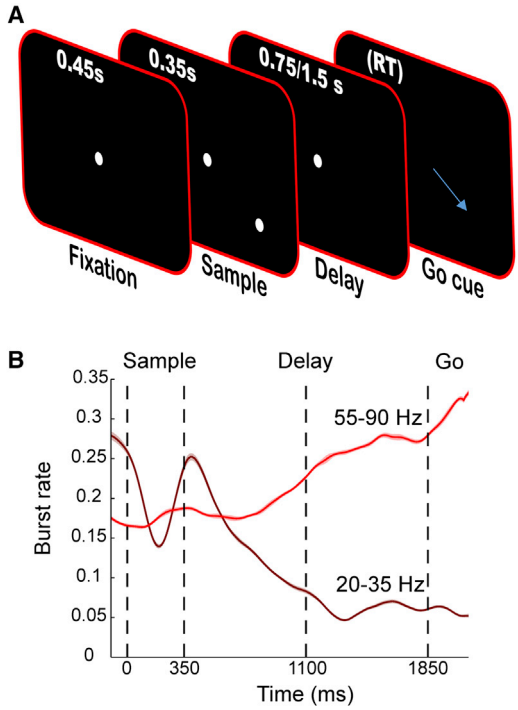

Figure 8. Bursts in a Delayed Saccade Task

(A) Schematic of the task. The monkeys were to maintain fixation until the fixation dot disappeared and then saccade to the earlier flashed target location.

(B) Burst rate in gamma (light) and beta (dark) frequency ranges averaged across monkeys and areas (FEF and PFC). Time 0 corresponds to the onset of the sample presentation. Displayed is the subset of trials (1/7) where the delay time was doubled (1,500 ms) relative to the standard delay of 750 ms (marked with dashed line 1,100 ms into trial). Shaded areas (hardly visible) represent the SE (n = 124).

the likely external stimulus. These non-linear activations are coupled with induced, stimulus-specific patterns of gamma and suppressed beta power (Lundqvist et al., 2010). It is the addition of synaptic potentiation (Wang et al., 2006) to this network that makes it capable of storing and replaying stimulus-induced attractor states and, thus, useful for WM (Mongillo et al., 2008; Lundqvist et al., 2011). Therefore, the beta and gamma dynamics observed here seem to be a reoccurring theme, reflecting similar mechanisms that have been adapted across the cortex.

EXPERIMENTAL PROCEDURES

Experimental Setup

Two monkeys (1 Macaca mullata [M1] and 1 Macaca fascicularis [M2], maintained in accordance with the National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee for Animal Care) were trained on a serial WM task. To obtain a liquid reward (apple juice), monkeys had to memorize a series of two or three colored squares. Each side of the square was one degree visual angle (DVA) in length and placed either 4 or 6 (for horizontal squares only) DVA lateral of the central fixation, either at the horizontal meridian or 75° above or below the meridian. Stimuli were only shown on the visual hemifield contralateral to the recording site. All behavioral procedures were controlled using MonkeyLogic, a MATLAB-based software tool for control of behavioral experiments (Asaad and Eskandar, 2008) interfacing through digital IO cards (NI-PCI6221, National Instruments). Eye movements were controlled using a video-based tracking system (Eyelink 1000, SR Research).

Data Collection

For each recording, a new set of acute electrode pairs (tungsten, epoxy-coated, FHC) was lowered through a grid (19-mm diameter, custom-designed grid and microdrives) that was always oriented in the same direction. A total of 18 (M1) and 20 (M2) prefrontal electrodes were recorded from simultaneously on each session (14 sessions for M1 and 16 session for M2), but only the ones containing isolatable units were kept for further analysis. The recordings were performed on a Blackrock Cerebus system and Plexon unity-gain headstages. The LFPs were recorded at a sampling rate of 30 kHz.

Signal Processing

We first removed apparent noise sources from the signal. In particular, a notch filter was applied to remove 60-Hz line noise with constant phase across a session. In addition, we removed periodic deflections seen in the evoked potentials (every 47 ms, lasting 1 ms, on a subset of electrodes, phase locked to stimulus onset). The signal was filtered and downsampled to 1 kHz. We applied and compared three different methods for time-dependent spectral estimation of the signal: Morlet wavelet analysis (Tallon-Baudry et al., 1996) as well as multi-taper analysis (with a family of orthogonal tapers produced by Slepian functions; Slepian, 1978; Thomson, 1982; Jarvis and Mitra, 2001) and band-pass filtering techniques (Butterworth filter of order 4–6 applied in the forward and reverse directions to obtain the effect of zero-phase filtering, followed by Hilbert transform to extract analytic representation of the signals). They all yielded very similar results in terms of qualitative time-frequency content. They also led to comparable burst extraction outcomes. For all the presented spectrograms (except Figures S6 and S2 in which a large frequency range was scanned and Morlet wavelets used; number of waves = 6, number of octaves = 6, and a step size of 0.05 octaves) and for all burst extraction, the multi-taper approach was adopted with frequency-dependent window lengths corresponding to six to eight oscillatory cycles and frequency smoothing corresponding to 0.2–0.3 of the central frequency, f_0, i.e., $f_0 \pm 0.2f_0$, where f_0 were sampled with the resolution of 1 Hz (this configuration implies that two to three tapers were used). The spectrograms were estimated with the temporal resolution of 1 ms.

Typically we present total power of raw LFPs (after removal of noise) without subtracting any baseline or estimated evoked content. Evoked (Figure S2B) power, in contrast, was calculated by averaging the LFPs across trials for each electrode and estimating power on the averaged signal.

Selection of Gamma-Modulated Sites

Spectral power in correct trials was estimated for each electrode. Electrodes with obvious artifacts (n = 14) were removed. These artifacts manifested themselves as persistent narrow-band patterns within either 6–8 or 20–22 Hz (n = 6) and apparent bleeding of spikes into the upper frequency ranges of the spectral power estimations (above 50 Hz, n = 8 in M2). Each of the remaining 321 electrodes was categorized either as a so-called gamma-modulated or non-modulated site. Gamma-modulated sites were defined as those that displayed a significant increase in gamma-band power during all presentations for either two- or three-item trials (comparing the last 200 ms of the 300-ms presentation window to the 200-ms interval preceding the presentation; Wilcoxon signed-rank test at $\alpha = 0.05$).

Burst Extraction

In the first step of the oscillatory burst identification, a temporal profile of the LFP spectral content within a frequency band of interest was estimated. We used two alternative methods of spectral quantification (see above). We either narrow-band-filtered LFP trials and extracted the analytic amplitudes (envelope) or we used single-trial spectrograms, obtained with the multi-taper approach, to calculate smooth estimates of time-varying band power (all presented results were obtained with the multi-taper approach; the results for the two methods were very similar). Next we defined oscillatory bursts as intervals during individual trials when the respective measure of instantaneous spectral

power exceeded the threshold set as two SDs above the trial mean value for that particular frequency, and with the duration of at least three cycles. Further, we extracted the time-frequency representation of the signal in the spectro-temporal neighborhood of each burst using the multi-taper approach, and we averaged the power density across the local spectral range to obtain its local temporal profile. We resorted to fitting two-dimensional Gaussian function to the local time-frequency map to specify the aforementioned neighborhood. Finally, we defined the burst length as a time subinterval where the average instantaneous power was higher than half of the local maximum (half-power point).

Having the burst intervals extracted for the beta band (20–35 Hz) and three gamma sub-band oscillations (40–65, 55–90, and 70–100 Hz) from each trial, we defined a trial-average measure, a so-called burst rate for each spectral band. This quantity corresponds to the chance of a burst occurrence on an individual electrode at a particular time in the trial (a proportion of trials where a given electrode displays burst-like oscillatory dynamics around the time point of interest sliding over the trial length).

To estimate how narrow in frequency the burst events were, we defined their frequency span analogously as their temporal length, i.e., in terms of the half-power point (3-dB drop). Hence, similarly as in the burst length estimation, we averaged the time-frequency representation of the signal in the local neighborhood of the burst, this time along the temporal dimension, and we identified the frequency range where this average spectral power component did not fall below 50% of the local maximum (epicenter of the burst's power).

Burst Auto-correlation and Cross-Correlation

To describe the temporal characteristics of the burst occurrence, we estimated the auto-correlation and cross-correlation functions. We first extracted a binary string for each electrode in each trial where only the middle time point of every burst was set to 1. Then for all electrodes we evaluated the trial-average auto-correlation using this binary representation. Analogously, we averaged cross-correlation between the binary representations of burst occurrences at all pairs of electrodes. In the case of both auto- and cross-correlations, the estimates were corrected for the lag-dependent size of overlap between the two binary strings, i.e., raw correlations were scaled by the length of the strings minus the size of lag.

We also estimated the CV_2 for inter-burst intervals (Holt et al., 1996) using the binary string representation of burst occurrences at each electrode according to the following formula:

$$CV_2 = \frac{1}{n-1}\sum_{i=1}^{n-1}\frac{2|I_i - I_{i+1}|}{I_i + I_{i+1}}$$

where I_i is the time interval between the ith consecutive pair of burst events (middle time points) (i = 1, .., n) across the concatenated binary strings of all trials.

Burst-Spike Interactions

To calculate the trial-average firing rates within and outside bursts, the instantaneous firing rates were first estimated for each unit by convolving spike trains with a Gaussian kernel (20 ms wide). Then, with a sliding window, we averaged normalized firing rates across trials where the matching electrode was either engaged in an oscillatory burst (within-burst firing) or it was not (outside-burst firing). The aforementioned normalization of firing rates was performed independently for each trial and unit by calculating the relative firing rate change with respect to the trial mean value (within-trial modulation).

Selection of Informative Cells

The bias-corrected PEV (Olejnik and Algina, 2003; Buschman et al., 2011) was estimated from average firing rates in 50-ms bins across trials with different stimulus-dependent conditions. We performed two-way ANOVA where trials were either grouped by the location (Figure S4 only) or combination of location and color of the presented items. All correct trials were used, as the groups were well balanced each session. The bias correction was used as it avoids the problem of non-zero mean PEV (ω^2) for small sample sizes.

$$\omega^2 = \frac{SS_{Between\ groups} - df \times MSE}{SS_{Total} + MSE},$$

where MSE is the mean squared error, df is the degrees of freedom, SS_{Total} is the total variance, and $SS_{Between\ groups}$ is the variance between groups. As a result, (bias-corrected) PEV allowed for the quantification of information carried by the modulation of firing rates of individual units accounting for the location or combination of color and position of the presented stimulus object. A unit was defined as informative if its PEV for the combination of color and location exceeded the threshold of 0.05 in two consecutive time bins anywhere in the trial after the relevant stimulus had been presented. Normalization of firing rates for each cell amounted to the scaling by its total intra- (across time) and inter-trial average firing rate.

Statistical Methods

The majority of tests performed in this study were nonparametric due to insufficient evidence for model data distributions. To address the multi-comparisons problem, we employed Kruskal-Wallis, Friedman's, and Wilcoxon's signed-rank tests where appropriate. In addition, for the comparison between temporal profiles of the normalized firing rates within versus outside oscillatory bursts, we resorted to a permutation test on the largest cluster-based statistics (Maris and Oostenveld, 2007), originally proposed to increase the test sensitivity based on the known properties of the data (here being temporal dependency). Finally, some attention should be given to the way we report correlations between the measures of time-varying spectral band content and burst rate statistics. The correlation analyses were performed on individual electrodes and only the summary statistics (mean and SE) for the electrode-wise significant effects ($p < 0.01$) are presented.

Additional Analysis

We also performed spike-field coupling, cross-frequency coupling, and behavioral analyses, which were omitted from the main paper. The motivation and more details are described in the Supplemental Experimental Procedures.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures and eight figures and can be found with this article online at http://dx.doi.org/10.1016/j.neuron.2016.02.028.

AUTHOR CONTRIBUTIONS

Conceptualization, E.K.M., J.R., M.L., P.H., and T.J.B; Investigation, J.R. (main), T.J.B., M.L., and S.L.B.; Formal Analysis, M.L. and P.H.; Validation, P.H.; Data Curation, J.R., S.L.B., M.L., and P.H.; Software, P.H. and M.L.; Writing – Original Draft, M.L. and E.K.M.; Writing – Review & Editing, M. L., E.K.M., P.H., and J.R.; Visualization, M.L., P.H., and J.R.; Funding Acquisition, E.K.M., J.R., M.L., and P.H.; Resources, E.K.M.; Supervision, E.K.M.

ACKNOWLEDGMENTS

We thank Jefferson Roy, Simon J. Kornblith, Jacob Donoghue, Matthew Silver, Andre Bastos, Robert Schmidt, Roman Loonis, and Alik Widge for valuable suggestions and comments. This work was supported by the National Institute of Mental Health (NIMH, 5R01MH091174-05 and 5R37MH087027-07), The MIT Picower Innovation Fund, Human Brain Project (to P.H.), Junior Faculty Development Program (JFDP) Fellowship (to M.L.), and Deutsche Forschungsgemeinschaft (DFG) RO 3864/2-1 (to J.R.).

Received: July 17, 2015
Revised: December 22, 2015
Accepted: February 10, 2016
Published: March 17, 2016

REFERENCES

Amit, D.J., and Brunel, N. (1997). Model of global spontaneous activity and local structured activity during delay periods in the cerebral cortex. Cereb. Cortex 7, 237–252.

Asaad, W.F., and Eskandar, E.N. (2008). A flexible software tool for temporally-precise behavioral control in Matlab. J. Neurosci. Methods *174*, 245–258.

Axmacher, N., Henseler, M.M., Jensen, O., Weinreich, I., Elger, C.E., and Fell, J. (2010). Cross-frequency coupling supports multi-item working memory in the human hippocampus. Proc. Natl. Acad. Sci. USA *107*, 3228–3233.

Brunel, N., and Wang, X.J. (2003). What determines the frequency of fast network oscillations with irregular neural discharges? I. Synaptic dynamics and excitation-inhibition balance. J. Neurophysiol. *90*, 415–430.

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

Compte, A., Brunel, N., Goldman-Rakic, P.S., and Wang, X.J. (2000). Synaptic mechanisms and network dynamics underlying spatial working memory in a cortical network model. Cereb. Cortex *10*, 910–923.

Cromer, J.A., Roy, J.E., and Miller, E.K. (2010). Representation of multiple, independent categories in the primate prefrontal cortex. Neuron *66*, 796–807.

Donoghue, J.P., Sanes, J.N., Hatsopoulos, N.G., and Gaál, G. (1998). Neural discharge and local field potential oscillations in primate motor cortex during voluntary movements. J. Neurophysiol. *79*, 159–173.

Durstewitz, D., and Seamans, J.K. (2006). Beyond bistability: biophysics and temporal dynamics of working memory. Neuroscience *139*, 119–133.

Fisch, L., Privman, E., Ramot, M., Harel, M., Nir, Y., Kipervasser, S., Andelman, F., Neufeld, M.Y., Kramer, U., Fried, I., and Malach, R. (2009). Neural "ignition": enhanced activation linked to perceptual awareness in human ventral stream visual cortex. Neuron *64*, 562–574.

Fries, P. (2009). Neuronal gamma-band synchronization as a fundamental process in cortical computation. Annu. Rev. Neurosci. *32*, 209–224.

Fuentemilla, L., Penny, W.D., Cashdollar, N., Bunzeck, N., and Düzel, E. (2010). Theta-coupled periodic replay in working memory. Curr. Biol. *20*, 606–612.

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.

Fuster, J.M., and Alexander, G.E. (1971). Neuron activity related to short-term memory. Science *173*, 652–654.

Goldman-Rakic, P.S. (1995). Cellular basis of working memory. Neuron *14*, 477–485.

Gray, C.M., and Singer, W. (1989). Stimulus-specific neuronal oscillations in orientation columns of cat visual cortex. Proc. Natl. Acad. Sci. USA *86*, 1698–1702.

Gregoriou, G.G., Gotts, S.J., Zhou, H., and Desimone, R. (2009). High-frequency, long-range coupling between prefrontal and visual cortex during attention. Science *324*, 1207–1210.

Holt, G.R., Softky, W.R., Koch, C., and Douglas, R.J. (1996). Comparison of discharge variability in vitro and in vivo in cat visual cortex neurons. J. Neurophysiol. *75*, 1806–1814.

Honkanen, R., Rouhinen, S., Wang, S.H., Palva, J.M., and Palva, S. (2015). Gamma oscillations underlie the maintenance of feature-specific information and the contents of visual working memory. Cereb. Cortex *25*, 3788–3801.

Howard, M.W., Rizzuto, D.S., Caplan, J.B., Madsen, J.R., Lisman, J., Aschenbrenner-Scheibe, R., Schulze-Bonhage, A., and Kahana, M.J. (2003). Gamma oscillations correlate with working memory load in humans. Cereb. Cortex *13*, 1369–1374.

Jarvis, M.R., and Mitra, P.P. (2001). Sampling properties of the spectrum and coherency of sequences of action potentials. Neural Comput. *13*, 717–749.

Jensen, O., and Tesche, C.D. (2002). Frontal theta activity in humans increases with memory load in a working memory task. Eur. J. Neurosci. *15*, 1395–1399.

Kornblith, S., Buschman, T.J., and Miller, E.K. (2015). Stimulus load and oscillatory activity in higher cortex. Cereb. Cortex, bhv182.

Kucewicz, M.T., Cimbalnik, J., Matsumoto, J.Y., Brinkmann, B.H., Bower, M.R., Vasoli, V., Sulc, V., Meyer, F., Marsh, W.R., Stead, S.M., and Worrell, G.A. (2014). High frequency oscillations are associated with cognitive processing in human recognition memory. Brain *137*, 2231–2244.

Lachaux, J.P., George, N., Tallon-Baudry, C., Martinerie, J., Hugueville, L., Minotti, L., Kahane, P., and Renault, B. (2005). The many faces of the gamma band response to complex visual stimuli. Neuroimage *25*, 491–501.

Lisman, J.E., and Idiart, M.A. (1995). Storage of 7 +/- 2 short-term memories in oscillatory subcycles. Science *267*, 1512–1515.

Lundqvist, M., Rehn, M., Djurfeldt, M., and Lansner, A. (2006). Attractor dynamics in a modular network model of neocortex. Network *17*, 253–276.

Lundqvist, M., Compte, A., and Lansner, A. (2010). Bistable, irregular firing and population oscillations in a modular attractor memory network. PLoS Comput. Biol. *6*, e1000803.

Lundqvist, M., Herman, P., and Lansner, A. (2011). Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. J. Cogn. Neurosci. *23*, 3008–3020.

Lundqvist, M., Herman, P., and Lansner, A. (2012). Variability of spike firing during θ-coupled replay of memories in a simulated attractor network. Brain Res. *1434*, 152–161.

Maris, E., and Oostenveld, R. (2007). Nonparametric statistical testing of EEG- and MEG-data. J. Neurosci. Methods *164*, 177–190.

Miller, E.K., and Cohen, J.D. (2001). An integrative theory of prefrontal cortex function. Annu. Rev. Neurosci. *24*, 167–202.

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.

Mongillo, G., Barak, O., and Tsodyks, M. (2008). Synaptic theory of working memory. Science *319*, 1543–1546.

Nir, Y., Fisch, L., Mukamel, R., Gelbard-Sagiv, H., Arieli, A., Fried, I., and Malach, R. (2007). Coupling between neuronal firing rate, gamma LFP, and BOLD fMRI is related to interneuronal correlations. Curr. Biol. *17*, 1275–1285.

Olejnik, S., and Algina, J. (2003). Generalized eta and omega squared statistics: measures of effect size for some common research designs. Psychol. Methods *8*, 434–447.

Palva, J.M., Palva, S., and Kaila, K. (2005). Phase synchrony among neuronal oscillations in the human cortex. J. Neurosci. *25*, 3962–3972.

Pasternak, T., and Greenlee, M.W. (2005). Working memory in primate sensory systems. Nat. Rev. Neurosci. *6*, 97–107.

Pesaran, B., Pezaris, J.S., Sahani, M., Mitra, P.P., and Andersen, R.A. (2002). Temporal structure in neuronal activity during working memory in macaque parietal cortex. Nat. Neurosci. *5*, 805–811.

Rainer, G., and Miller, E.K. (2002). Timecourse of object-related neural activity in the primate prefrontal cortex during a short-term memory task. Eur. J. Neurosci. *15*, 1244–1254.

Ray, S., and Maunsell, J.H. (2011). Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. PLoS Biol. *9*, e1000610.

Roux, F., and Uhlhaas, P.J. (2014). Working memory and neural oscillations: α-γ versus θ-γ codes for distinct WM information? Trends Cogn. Sci. *18*, 16–25.

Roux, F., Wibral, M., Mohr, H.M., Singer, W., and Uhlhaas, P.J. (2012). Gamma-band activity in human prefrontal cortex codes for the number of relevant items maintained in working memory. J. Neurosci. *32*, 12411–12420.

Sandberg, A., Tegnér, J., and Lansner, A. (2003). A working memory model based on fast Hebbian learning. Network *14*, 789–802.

Sederberg, P.B., Kahana, M.J., Howard, M.W., Donner, E.J., and Madsen, J.R. (2003). Theta and gamma oscillations during encoding predict subsequent recall. J. Neurosci. *23*, 10809–10814.

Shafi, M., Zhou, Y., Quintana, J., Chow, C., Fuster, J., and Bodner, M. (2007). Variability in neuronal activity in primate cortex during working memory tasks. Neuroscience *146*, 1082–1108.

Slepian, D. (1978). Prolate spheroidal wave functions, Fourier analysis, and uncertainty. V-The discrete case. Bell Syst. Tech. J. *57*, 1371–1430.

Stokes, M.G. (2015). 'Activity-silent' working memory in prefrontal cortex: a dynamic coding framework. Trends Cogn. Sci. *19*, 394–405.

Stokes, M.G., Kusunoki, M., Sigala, N., Nili, H., Gaffan, D., and Duncan, J. (2013). Dynamic coding for cognitive control in prefrontal cortex. Neuron *78*, 364–375.

Tallon-Baudry, C., and Bertrand, O. (1999). Oscillatory gamma activity in humans and its role in object representation. Trends Cogn. Sci. *3*, 151–162.

Tallon-Baudry, C., Bertrand, O., Delpuech, C., and Pernier, J. (1996). Stimulus specificity of phase-locked and non-phase-locked 40 Hz visual responses in human. J. Neurosci. *16*, 4240–4249.

Thomson, D.J. (1982). Spectrum estimation and harmonic analysis. Proc. IEEE Inst. Electr. Electron. Eng. *70*, 1055–1096.

Wang, Y., Markram, H., Goodman, P.H., Berger, T.K., Ma, J., and Goldman-Rakic, P.S. (2006). Heterogeneity in the pyramidal network of the medial prefrontal cortex. Nat. Neurosci. *9*, 534–542.

Watrous, A.J., Tandon, N., Conner, C.R., Pieters, T., and Ekstrom, A.D. (2013). Frequency-specific network connectivity increases underlie accurate spatiotemporal memory retrieval. Nat. Neurosci. *16*, 349–356.