






Working memory control dynamics follow principles of spatial computing

Received: 16 March 2022

Accepted: 7 February 2023

Published online: 14 March 2023

 Check for updates

Mikael Lundqvist ^{1,2}✉, Scott L. Brincat², Jonas Rose ^{2,3},
Melissa R. Warden ^{2,4}, Timothy J. Buschman ^{2,5}, Earl K. Miller^{2,7} &
Pawel Herman ^{6,7}✉

Working memory (WM) allows us to remember and selectively control a limited set of items. Neural evidence suggests it is achieved by interactions between bursts of beta and gamma oscillations. However, it is not clear how oscillations, reflecting coherent activity of millions of neurons, can selectively control individual WM items. Here we propose the novel concept of spatial computing where beta and gamma interactions cause item-specific activity to flow spatially across the network during a task. This way, control-related information such as item order is stored in the spatial activity independent of the detailed recurrent connectivity supporting the item-specific activity itself. The spatial flow is in turn reflected in low-dimensional activity shared by many neurons. We verify these predictions by analyzing local field potentials and neuronal spiking. We hypothesize that spatial computing can facilitate generalization and zero-shot learning by utilizing spatial component as an additional information encoding dimension.

Working memory (WM) is a mental sketchpad for the short-term storage and top-down control of information^{1–8}. This control is key to WM's central role in cognition⁹. We can select what we retain, read out or delete from WM as well as manipulate the contents^{7,8,10–17}. It is not clear however what neural mechanisms underlie such flexible control.

Recent work on the temporal neural dynamics of WM has begun to provide some insight. The central idea is that top-down control stems from interactions between bursts of gamma and beta power^{11,17–19}. The gamma bursts are associated with spiking that in turn encodes and maintains WM content. Beta bursts act as the control signal. They carry top-down information and inhibit gamma/spiking, thus controlling the access to WM contents carried by the spiking. This is supported by empirical observations of an anti-correlated “push-pull” relationship between beta and gamma during the encoding, read-out, and deletion of the contents of WM¹¹. For example, when information is encoded into, or read out from, WM, beta decreases and

gamma increases along with spiking carrying the WM content. When the content is no longer relevant and could be cleared out of the WM, the opposite dynamics is seen¹¹.

However, a key question remains: How can these gamma-beta interactions, which reflect the combined activity of millions of neurons, be selective enough to control the contents of individual items in WM? WM control, after all, is more than just turning WM “on” and “off”. It also includes selective operations on the individual items held in WM.

Here we propose a novel principle we dub *spatial computing* where selective control comes from utilizing network space. Spatial computing rests on the assumption that WM item representations are consistently moved across the spatial dimensions of the cortical network depending on task demands. Assigning or moving the representation of a WM item from one part of the network to another is considered a computation in itself. It is used to assign or change the status of an item, for instance its temporal order, if it is a sample or cue

¹Division of Psychology, Department of Clinical Neuroscience, Karolinska Institute, Stockholm, Sweden. ²The Picower Institute for Learning and Memory, Department of Brain and Cognitive Sciences, Massachusetts Institute of Technology, 43 Vassar Street, Cambridge, MA 02139, USA. ³Faculty of Psychology, Neural Basis of Learning, Ruhr University Bochum, 44801 Bochum, Germany. ⁴Department of Neurobiology and Behavior, Cornell University, Ithaca, NY 14853, USA. ⁵Princeton Neuroscience Institute, Princeton University, Washington Rd., Princeton, NJ 08540, USA. ⁶Division of Computational Science and Technology, School of Electrical Engineering and Computer Science & Digital Futures, KTH Royal Institute of Technology, Stockholm 100 44, Sweden. ⁷These authors jointly supervised this work: Earl K. Miller, Pawel Herman. ✉e-mail: mikael.lundqvist@ki.se; paherman@kth.se

item, or if it is currently prioritized or not. Control thus comes from *where* in network space a specific WM item is held. This allows the item to be accessed and operated on just by knowing its place in network space. Importantly, it enables control without having to know the precise network connectivity forming the ensemble for that item. In this view, spatial computing is mediated by a low-dimensional (where neurons share similar activity profiles) pattern of gamma-beta power across a network. The control operations on WM contents are reflected in the shared, low-dimensional components of neural activity²⁰. Spiking carrying the items held in WM is a high-dimensional (where neurons have independent activity profiles) component that appears where gamma is high and beta is low at each moment in time. Applying a set of WM operations (e.g., executing task rules) corresponds to imposing a low-dimensional spatio-temporal pattern on the network.

To test this idea we compared beta, gamma and spiking activity in several WM tasks that required several aspects of WM control. Spatial computing predicts that these control operations can be read from a low-dimensional spatio-temporal pattern of gamma and beta since this information is spatially organized. The control operations include the ordering of the objects and assigning their different uses (encoding/maintenance vs determining whether they match a test object). By contrast, gamma and beta are not expected to carry the identity of the items due to their coarser spatial scale than individual neurons. This information should instead be reflected in the high-dimensional spiking because it arises from the connectivity of single neurons, organized on a much finer spatial scale. Further, because the low-dimensional activity reflects WM operations not content, it should have a stable spatial pattern across different sessions where different sets of objects are used. In spiking, we should thus observe a mix of high-dimensional activity reflecting item identity and the low-dimensional components that account for the current operation being performed on them.

We found support for these predictions and thus offer a novel, spatial perspective on low-dimensional activity and the role of oscillations in WM control. We discuss how spatial computing offers a new perspective on mixed selectivity, neural subspaces and redundancy in cortex. Importantly, we discuss how spatial computing provides an account for the generalizability of WM. In contrast to many computational models of WM, spatial computing allows items novel to a specific task to be operated on without having to re-train networks to the new items.

Results

We analyzed multiple-electrode neurophysiological recordings from the prefrontal cortex (PFC) of five rhesus monkeys performing four WM tasks. In three of the tasks the monkeys had to remember sequences of objects or colored squares. Two tasks were single-object delayed match-to-sample tasks.

We tested the following predictions of spatial computing: 1. There are different neural sources for control-related activity (i.e., top-down task demands) vs the specific items held in WM. 2. The control-related activity is organized as spatial patterns. 3. The patterns are stable across different recording sessions that use different sets of WM items but have the same top-down demands. 4. The control-related activity modulates the spiking of individual neurons such that their activity reflects both the identity of the item held in WM and current task demands.

Prediction 1: Independent neural sources of control-related and item-specific WM activity

To test this prediction, we used local field potentials (LFPs) and single neuron activity recorded with acute electrodes in PFC of rhesus monkeys performing a sequence two-item WM task. The monkeys had to remember the identity and order of two visual objects presented in sequence (Task 1, Fig. 1a). After a brief memory delay, there was a two-object test sequence that could either be identical to the encoded

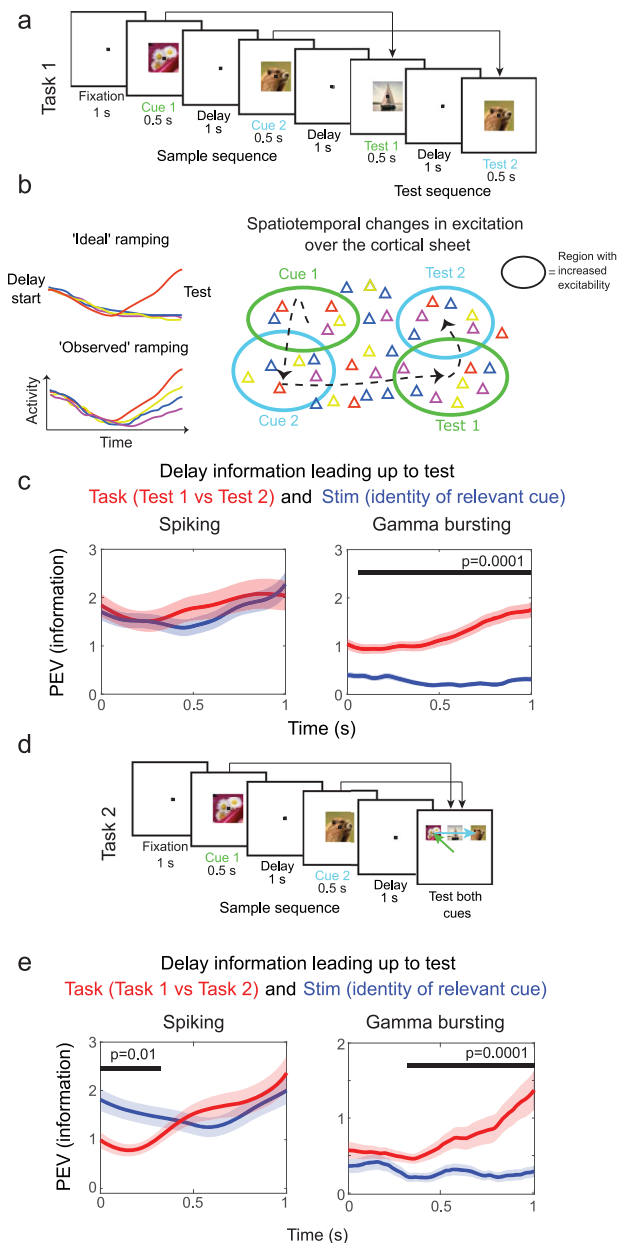
sample sequence or a mismatch (where either the temporal order or the identity of the objects was changed). After the second test object, monkeys indicated whether the test object sequence matched that of the sample object sequence seen at the start of the trial.

Both item-specific activity (carrying information about each object's identity) as well as control-related activity (i.e., the requirement to remember the order and to determine if they match a test sequence) are needed to solve the task (Fig. 1a¹³). Our previous work demonstrated a ramp-up of spiking (and gamma bursting) at the end of a memory delay that reflected a "read-out" of an item from WM^{11,19}. On a population level the ramp up in spiking was selective to order information. Out of the two items held in WM, spiking only increased for the item that was relevant for the upcoming test (information about Sample 1 item before Test 1, information about Sample 2 item before Test 2 in Fig. 1^{11,13}).

This selective ramp-up we observed on the population level was reflected in item-selective single neurons that only ramped up to one of the two tests (demonstrating order specificity). However, when individual neurons ramped before a specific test (either Test 1 or Test 2), they did so regardless if their preferred item had been presented during the corresponding sample or not (conceptually illustrated in Fig. 1b, single neuron examples shown in Fig. S1). An idealized, item-specific ramping neuron would selectively ramp up spiking only for its preferred stimulus, not for others (Fig. 1b, left). Spatial computing suggests that this effect is due to independent origins of control-related and item-specific activity. The ramping comes from top-down control-related excitation selective to the order of the two tests. It targets cortical locations at much coarser spatial scales than individual neurons. As a result, millions of neurons in these cortical locations receive excitation regardless of their item preference (Fig. 1b right). Due to the spatial integration of LFP activity over a few hundred micrometers, item-specific activity should largely cancel out in gamma (and beta) bursting. Therefore, spatial computing predicts that these LFP oscillations should capture spatially organized information about the order of items but not their identities per se. Importantly, spiking, which probes network activity at much finer spatial scales, should account for both item-specific activity arising from recurrent connectivity and control-related activity inherited from the gamma and beta interactions.

We thus examined spiking and gamma bursting in the delay periods of 1000 ms leading up to the first and second test stimuli (Fig. 1a). The data in these periods were labelled by the control-related context (Test 1 or Test 2) and by item identity (i.e., the item that was expected based on the sequence held in WM - the first item from sample sequence for Test 1 and the second item for Test 2). We then calculated the percentage of variance (PEV) in the neural activity explained by these labels. As predicted, the ramp-up of gamma bursting carried information mainly about the control component, order, i.e., whether it was the lead up to Test 1 or Test 2 cue, and not the identity of the retained items (Fig. 1c, right panel). Spiking carried a mixture of the two information components (Fig. 1c, left), i.e., both the order and the identity of the expected item. We also found a similar difference between spiking and beta activity in line with the prediction about the beta as a correlate of control signals (Fig. S2). As expected from the control signal, beta also mainly carried order information.

We further tested this prediction by analyzing data from another WM task (Fig. 1a, d). In Task 2, two cues were again sequentially presented, just as in Task 1, but then they differed in how those memories were tested after the memory delay. In Task 2, rather than making a yes/no decision about whether a test sequence matched the remembered one, three objects were presented simultaneously. The monkey had to choose, using eye movements, the correct two objects in the correct order. Monkeys switched back and forth between these two tasks. Note that both tasks are identical up until the testing period after the memory delay. Nonetheless, a previous analysis of this data showed that there were different patterns of spiking ramp-up at the end of the memory



delay reflecting the different task demands (Warden and Miller, 2010). Here, we found that control-related information (i.e. whether the monkey was performing Task 1 or Task 2) could be determined from spatial patterns of gamma bursting (Fig. 1e, right panel). As predicted by spatial computing, the ramp-up in spiking instead reflected not only the control-related information but also item identity (Fig. 1e, left).

We further elaborated on these results by using demixed principal component analysis (dPCA) of the two tasks (Fig. 2). dPCA decomposes population activity into principal components dependent on task parameters specified in the analysis²⁰. In our case, these relevant parameters modulating population activity were WM items used in the task, the task being performed and time (task epochs). Here we analyzed both gamma bursting and spiking activity using dPCA. We identified low-dimensional components, attributed to WM control in the spatial computing framework, using activity that varied across tasks or task epochs but did not depend on specific WM items. We also extracted high-dimensional components that varied with the specific items used in a trial. The control-related components captured two types of activity: differences between the two tasks (“Task individual”) and shared patterns of activity over time in the two tasks (“Task

Fig. 1 | Disassociation between gamma and spiking. **a** In Task 1, two object cues are presented and then tested sequentially. In “match trials”, to which monkeys had to respond, the order and identity of the objects have to be the same for sample and test sequences. **b** Cartoons to illustrate spatial computing (for actual data see Fig. S1). Left: Item-specific activity ramps up just before the information about item identity is needed in the task. In the idealized case of stimulus selective neurons, their activity only ramps up for one specific cued item (top). However, recorded neurons with ramping activity respond to all cued items to varying levels (bottom). Right: the model explains the observed phenomenon of post stimulus ramping after all stimuli to varying degrees by external time-varying excitation that travels over the cortical sheet. Neurons preferring for instance the red or blue cues are scattered throughout the network. Some sites exhibit ramping excitation, reflected in increased activity of all neurons at that site, though to the lesser degree when the preferred object was not provided. **c** Red plots show PEV accounting for test order effects estimated over two groups of trial periods, delay periods preceding either Test 1 or Test 2. Blue curves reflect PEV wrt. cued item identity (information about identity of Sample 1 prior to Test 1 and Sample 2 leading up to Test 2 stimuli). Black bars demonstrate when blue and red plots differ, using cluster based statistics ($n = 199$ for gamma, $n = 283$ for single neurons). **d** Task 2 structure, which is identical to Task 1 until Test 1. Unlike in Task 1 however, at Test 1 both Sample 1 and Sample 2 items are tested in parallel by monkeys making eye-movements to the two targets. **e** PEV information pooled over Task 1 and 2 in the delay period leading up to Test 1 ($n = 199$ for gamma, $n = 283$ for single neurons). Red plots show PEV estimated over periods grouped based on task (Task 1 vs 2), whereas blue curves illustrate PEV information about the cued item identity (average of information about Sample 1 and Sample 2). Shaded regions indicate standard error of the mean in all panels. Source data are provided as Source Data file. Panels a and d were created with images from Melissa R. Warden, Earl K. Miller, The Representation of Multiple Objects in Prefrontal Neuronal Delay Activity, *Cerebral Cortex*, Volume 17, Issue suppl_1, September 2007, Pages i41–i50, <https://doi.org/10.1093/cercor/bhm070> with permission from Oxford University Press.

general”). Item components reflected the item-specific spiking. As predicted, the control-related components were more dominant in gamma bursting (Fig. 2, top left) whereas spiking additionally captured strong item dependent components (Fig. 2, top right panel).

One possibility is that the observed difference between gamma and spiking simply reflected differences in the quality in these measures of cortical activity (spikes vs LFPs). In other words, item-specific information might simply be more difficult to read out from gamma. To test for this, we performed the same analysis on another WM dataset referred to as Task 3 (Fig. 3, Lundqvist et al. 2016). As in the two previous tasks, the items were presented sequentially. Though this time each sample was in a different extrafoveal location instead of all foveally as in Task 1 and Task 2. As a result, we expected the item-specific spiking to be spatially distributed due to spatial organization of location-related information in PFC^{21,22}. Consequently, item-specific activity should thus be strongly present also in gamma activity. Therefore we should not observe the same disassociation between control-related and item-specific activity in spiking and gamma, as in the previous tasks. Both gamma and spiking should now carry similar contents. The results shown in Fig. S3 are consistent with the prediction. There was similar amount of control-related and item-specific activity in both neural measures. This suggests that differences between the gamma and spiking reported in Tasks 1 and Task 2 reflected distinct spatial organization of control-related and item-specific information, and not just quality differences between gamma and spiking as neural measures. Further, in line with our earlier discussion of gamma-beta interactions as spatial computing correlates, we found that beta bursting was dominated by the low-dimensional, task dependent components as well (Fig. S3).

Predictions 2 and 3: Control-related activity form spatial patterns that are stable across sessions

Spatial computing predicts that control-related spiking and LFP activity is spatially distributed and can be decomposed into spatial

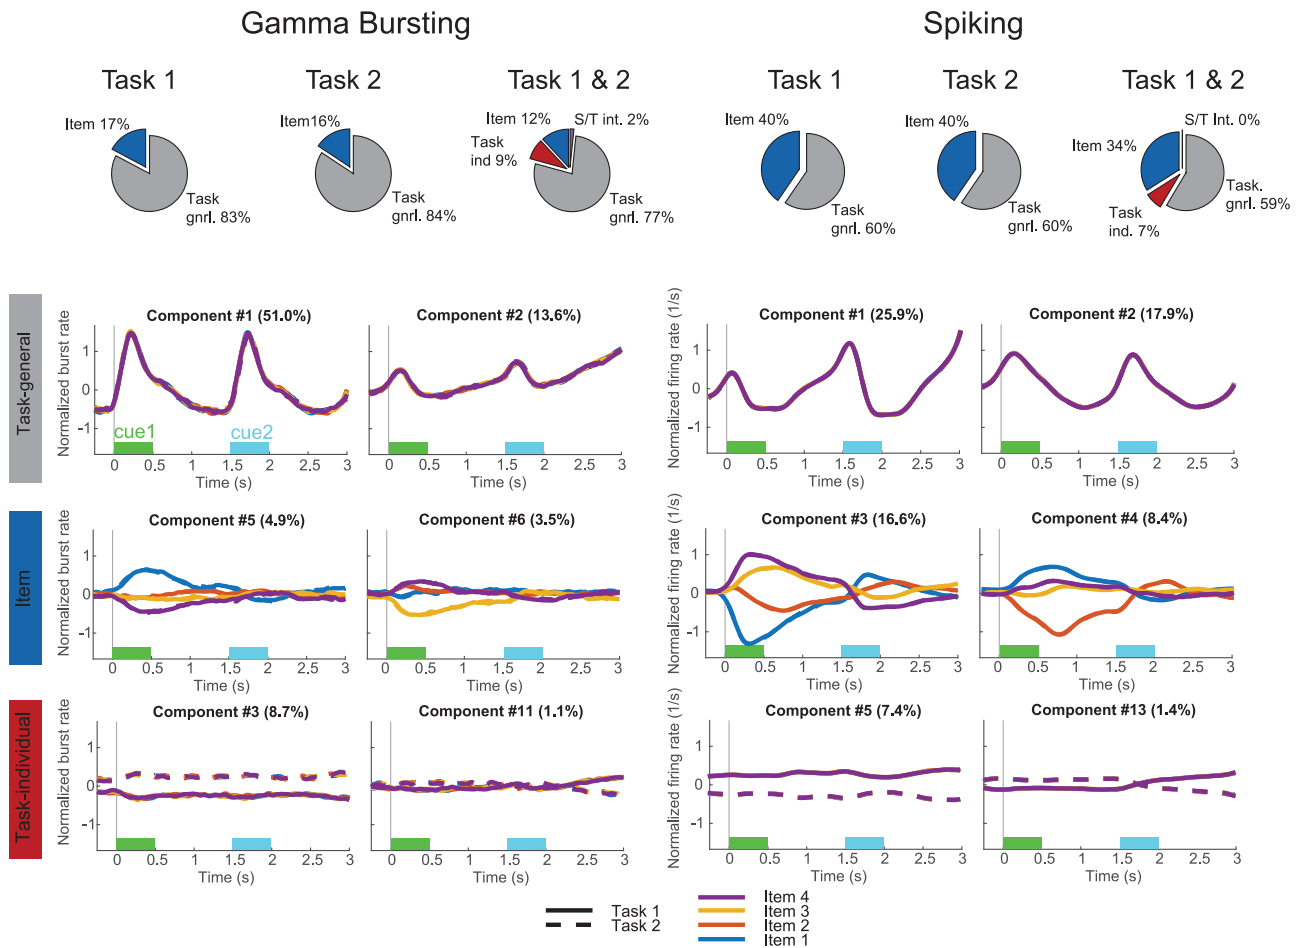


Fig. 2 | Demixed principal component analysis (dPCA) of gamma and spiking. Spatial computing implies that task dependent components (red and grey) should be more prevalent in gamma bursting as compared to spiking (c.f. Fig. 1). Here we used dPCA²⁰ to extract the principal components and attribute them to task control-related and item-specific activity ($n = 199$ for gamma, $n = 283$ for single neurons). “Task general” (grey) and “Task individual” (red) components correspond to low-dimensional task control-related activity. “Task general” components reflect shared patterns of activity over time in the two tasks whereas “Task individual” components explain the variance that originates from the differences between the two tasks. Item

dependent components (blue) account for the variance between four different cued items (item-specific activity). The green (sample 1) and light blue (sample 2) rectangles mark when the samples were shown. “S/T int” refers to components that depended both on item and task when data from both tasks were analyzed together. The bottom half of the figure shows example components for Task 1 and Task 2 combined. Shown are the first two components of each type (meaning several task general components are not shown as they tended to dominate). Source data are provided as Source Data file.

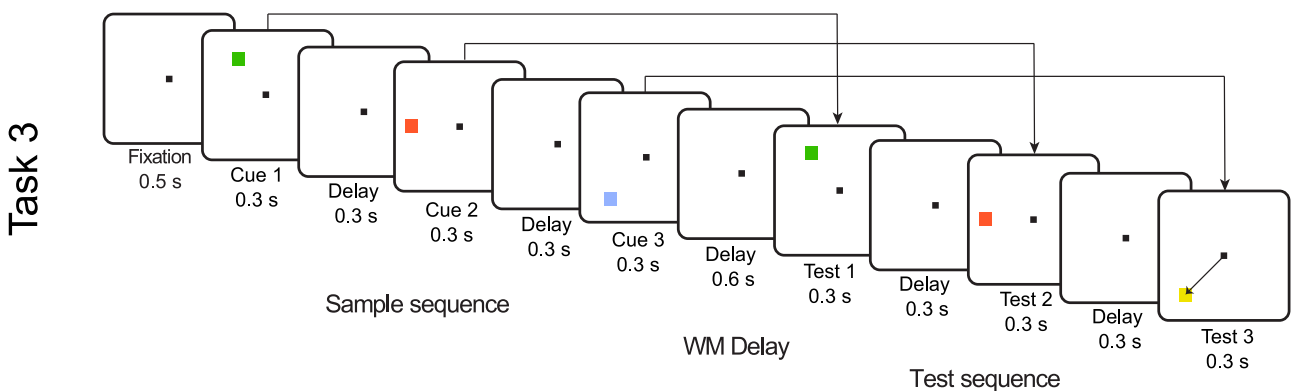


Fig. 3 | Task 3. In Task 3, two or three colored squares were presented in a sequence. Their location and color were to be remembered. Following the 0.6 s

delay, the sequence of squares was repeated. Monkeys needed to saccade to the square in the test sequence that changed color relative the sample sequence.

components reflecting patterns of excitation. To test this hypothesis, we used data from dense chronically implanted electrode arrays, which enabled us to 1) map the spatial distribution of various components, and 2) assess their stability across different recording sessions. We recorded simultaneously from 4 Utah arrays, implanted in left and right PFC. The monkey performed an object delayed match-to-sample task over multiple sessions (Task 4, Fig. 4). Different sets of 8 objects (WM items) were used for each session.

We focused on gamma activity since for each electrode it reflects the activity from a stable group of neurons in that spatial vicinity across sessions (whereas the same individual single neurons may not remain isolated across sessions). We performed dPCA analysis of the gamma activity (45–100 Hz) within the delay period in each session. We focused on the first two control-related components that reflected shared activity over time regardless of item identity. They explained on average 75 and 8% of the data variance, respectively. To localize the spatial origins of the two components in each recording session, we extracted their corresponding dPCA loadings (weights used for projecting the data). Further, to minimize the risk of overfitting the dPCA model to session specific noise we sparsified the weights. We only kept those weights that significantly contributed to the component (see Methods). We then averaged the dPCA weights across five sessions per component and used the two resulting weight vectors to derive the corresponding projections for each session. This yielded virtually identical components in each session (Fig. 4b, left). Further, session-specific dPCA components, i.e., based on loadings extracted from individual sessions, showed striking resemblance to those obtained with the five-session averaged weight vector (Fig. 4b, right). Thus the spatial distribution of these components was stable across sessions (and across sets of WM items). We plot their spatial organisation in Fig. 4c to illustrate that they were distributed in all four recorded areas and distinct for the two components. Thus, as predicted by spatial computing, different control-components corresponded to spatial patterns of activity.

The difference between the two control-related components is also reflected over time. The first component corresponded to activity that slowly decayed over the delay period. The second one corresponded to activity that initially decayed but then ramped up before the end of the delay. The components also exhibited differences during the sample period. The first component was elevated throughout the sample. The second component had activity primarily at the onset and offset of sample. Since the components were selected solely based on the delay activity, these differences in other task periods imply that they accounted for consistent variations between recording locations (and that they did not reflect overfitting to the data).

We interpreted the first component as controlling where sample information was encoded. The second component, with the ramp-up during the delay, determined the spatial location from where information was read out. We wanted to test more explicitly whether this ramp-up was related to control. In addition, we wanted to ensure that the distinct spatial distribution of activity across various task epochs did not just reflect the “passage of time” in the trials. Thus we recorded array data from yet another task, Task 5 (Fig. 5a). This task was similar to Task 4, but the delay duration varied randomly from 1 to 4 s. This made the timing of the upcoming test stimulus unpredictable. On 90% of trials, there was a pre-test (just prior to the test) cue alerting the monkey to the test. Behavioral analysis suggested that the monkey exploited pre-test cue as its accuracy was significantly higher on cue-trials (93.4% vs 87.1% correct, $p = 0.0001$, t -test with randomization). This allowed us to more explicitly link any neurophysiological components related to the pre-test cue with WM control of read-out from WM. The first two task-general dPCA components indeed had strong responses to the pre-test cue with elevated gamma activity following the cue (Fig. 5b). We used the dPCA coefficients extracted from trials with delays lasting 2 s to project data from the remaining trials with the 4 other delay lengths. This demonstrated that the elevated activity

towards the end of the delay did not merely reflect the passage of time. Instead it was in direct response to the up-coming test, consistently with a role in WM control. In addition, it meant the same spatial configuration of electrodes contributed to the read-out components regardless of the duration of the delays, consistently with the spatial distribution of activity being used for such WM control.

Interestingly, comparing the first two control-related coefficients from Task 4 with Task 5, they were only partially overlapping spatially (compare Fig. 5c with Fig. 4c). Projecting data from Task 4 using coefficients from Task 5 did indeed only partially recapture the components from the original analysis of Task 4 (Fig. S5). This suggested that despite the strong similarities between the two tasks, the spatial locations used for encoding and read-out were remapped across the tasks. Consequently, it implies that encoding or read-out locations in cortex are not completely hardwired. Instead they change dynamically with task demands, which may for instance reflect the need for more neural resources for encoding or read-out.

Prediction 4: Low-dimensional spatial excitation gives rise to context dependent selectivity

One of the implications of spatial computing is that at the single-neuron level there is a mix of control and item-specific activity. This arises out of control-related, low-dimensional spatial patterns of excitation combining with the item-selectivity of individual neurons. To test this hypothesis, we analysed spiking and LFPs from the third WM task (Task 3, Fig. 3). In this task the three samples were presented in unique locations so information about their order did not need to be stored. However, each sample and the corresponding test cue were presented in the same location and had to be distinguished somehow to facilitate their comparison. According to the principle of spatial computing, this should be achieved by having distinct spatial patterns of excitation at sample and test cues, respectively. It should then result in different sets of neurons encoding these two cues despite the fact that they were presented in the same location and that the stimulus preferences of individual neurons did not change (Fig. 6a). We considered here beta and gamma bursting as proxies for local excitation. In particular, beta bursts should reflect excitation levels distinguishable from spiking. This is because they express less cross-talk between spiking and LFP activity than gamma²³. We have previously shown that beta bursts are anti-correlated with bursts of gamma and spiking^{11,19}. For example, beta bursting is lower when the sample objects are presented while gamma bursting is then elevated (Fig. 6b; Fig. S6).

Here, we found that the levels of beta activity had a distinct shift downwards before and during test periods compared to sample encoding periods (Fig. 6b) suggesting a general shift in excitability. It was not due to saccades as it was seen equally during test periods in which animals responded or withheld the response (compare Test 1 (no response) and Test 2 (response) in Fig. 6b). This implies that excitation levels indexed by LFPs (and approximated here by beta) were generally elevated during test periods compared to encoding. In line with earlier analysis of the same data (Lundqvist et al. 2016), sites with neurons displaying item-selective activity had lower beta during encoding (sample cues) than sites with non-selective neurons (Fig. 6b; using units selective to Sample 1 in the first 600 ms after the sample onset). However, during the test periods beta for non-selective sites was consistently lower compared to the lowest beta observed at selective sites during encoding. This more widespread excitation (reflected in LFPs) should according to the hypothesis be manifested in a more widespread item-selectivity (of single neurons) in the network. In trials where the same items were presented during Sample 1 and Test 1 (to which the animals thus did not respond to), beta was suppressed (Fig. 6b), and gamma power (Figure S6) and firing rates were increased during the test relative to the sample period. Importantly, item-selective activity was confined to a smaller part of the network during the encoding compared to the test periods (Fig. 6c). A significantly larger portion of the

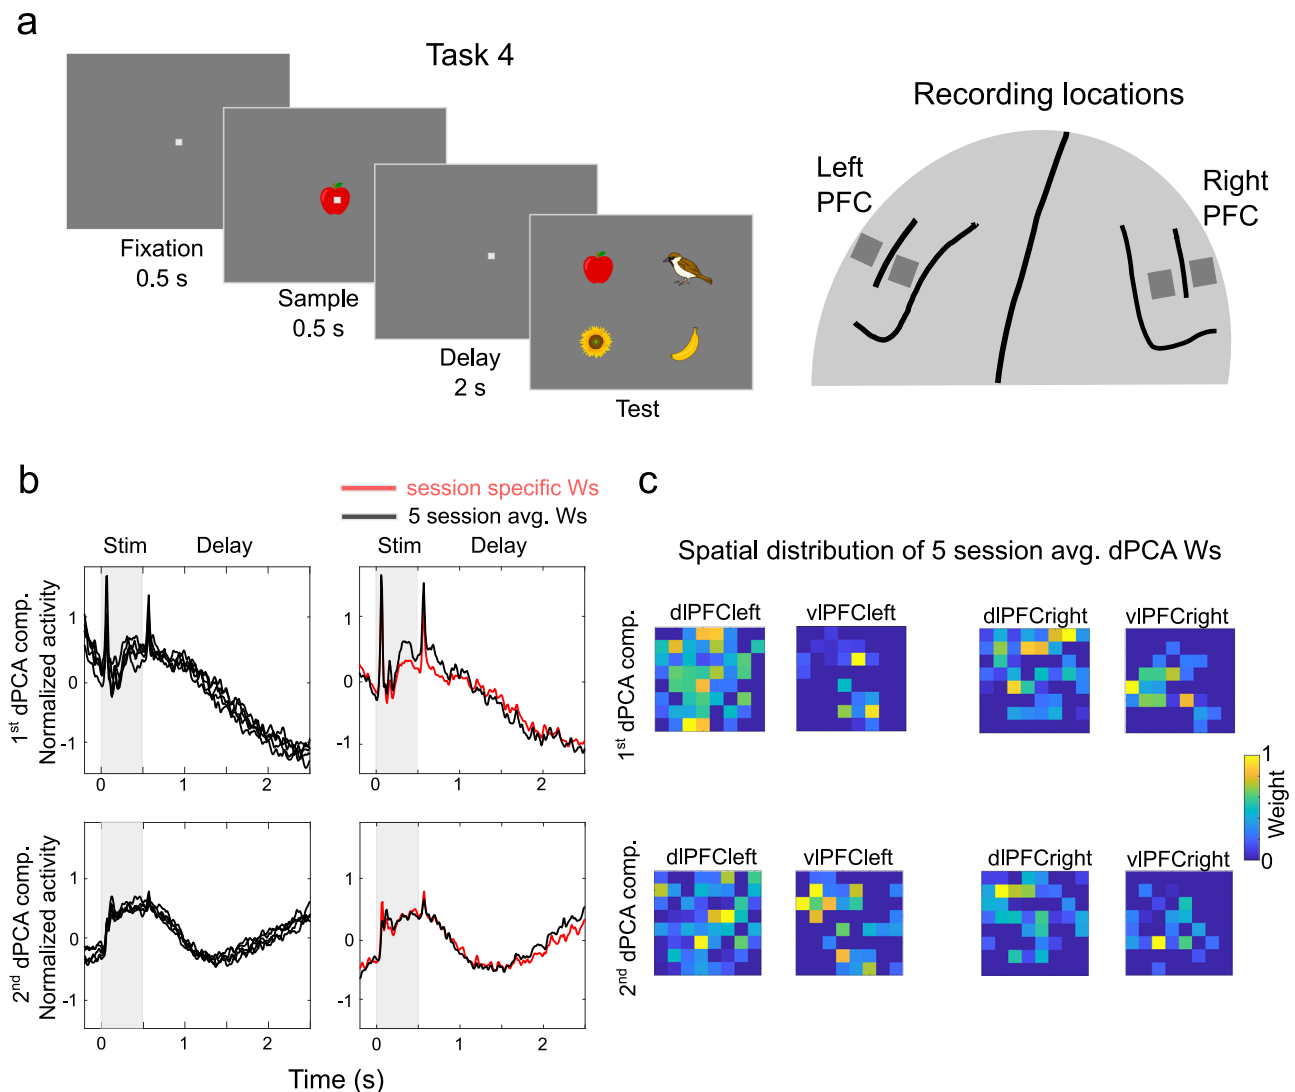


Fig. 4 | Spatial patterns of control-related activity in a delay-match-to-sample task. **a** Left: Schematic depiction of the delayed match-to-sample task. Right: schematic depiction of the recording locations of the four chronically implanted Utah arrays (left and right dIPFC, left and right vIPFC). **b** Left: the first (top) and second (bottom) task-general dPCA components of gamma activity projected using the five-session average dPCA weight vectors for each one of the five sessions. Right: plotted are the five-session averages of the first and second task-general

dPCA gamma components using i) the five-session average dPCA weight vectors (black) and ii) session-specific weight vectors (red). **c** The 4-array spatial distribution of the session-average normalized dPCA weights used in 'b' for the first (top) and second (bottom) task dependent dPCA components. It demonstrates that the different components formed unique and stable spatial patterns distributed over PFC. Source data are provided as Source Data file. Panel a was created with clip art images from © DESIGNALIKIE, Limited.

neurons responded to Test 1 than Sample 1 stimuli, even in trials in which they were identical (69/480 had significant variance explained by Sample 1 identity compared to 154/480 by Test 1 identity using the first 600 ms from the onset of sample cue and test cue, respectively; $p < 1e-11$, two-sided binomial test). This suggests that enforcing different spatio-temporal patterns of excitation during sample and test periods could be an effective way to distinguish information elicited in the selective neural population in the two task periods. This is a natural implication of the principle of spatial computing.

Discussion

We have tested the hypothesis that the PFC implements *spatial computing* for the control of WM. The central idea is that top-down control-related information (e.g., the ordering of multiple items held in WM, how they are used etc.) operates in network space. Different WM items are distributed across different locations in a network so they can be manipulated independently. The manipulations are reflected in neural oscillations that act on hundreds of thousands of neurons (rather than

target single neurons) at different network locations. The manifestation is a dynamic patchwork of beta and gamma bursts. As in our earlier framework, the gamma bursts co-register with the spiking carrying the individual items held in WM. Beta bursts act as an inhibitory control signal that inhibits gamma/spiking when and where beta power is high^{11,17,19}.

To understand how the spatial computing works consider our Task 1 requirement to retain two objects (A and B) in the order in which they appeared (first or second^{11,13}). This task, like many other WM tasks, requires several forms of control. First, input control where the sample objects are stored according to their order of presentation. Second, output control where only the currently relevant object is read out and compared to the first or second test objects. Third, the test objects also have to be presented as test objects rather than sample objects to facilitate the comparison. Here, in line with the principle of spatial computing, we found evidence that all these forms of WM control were aided by utilizing the spatial dimensions of the cortical sheet. The idea is that during the initial encoding of the sequence, different spatio-

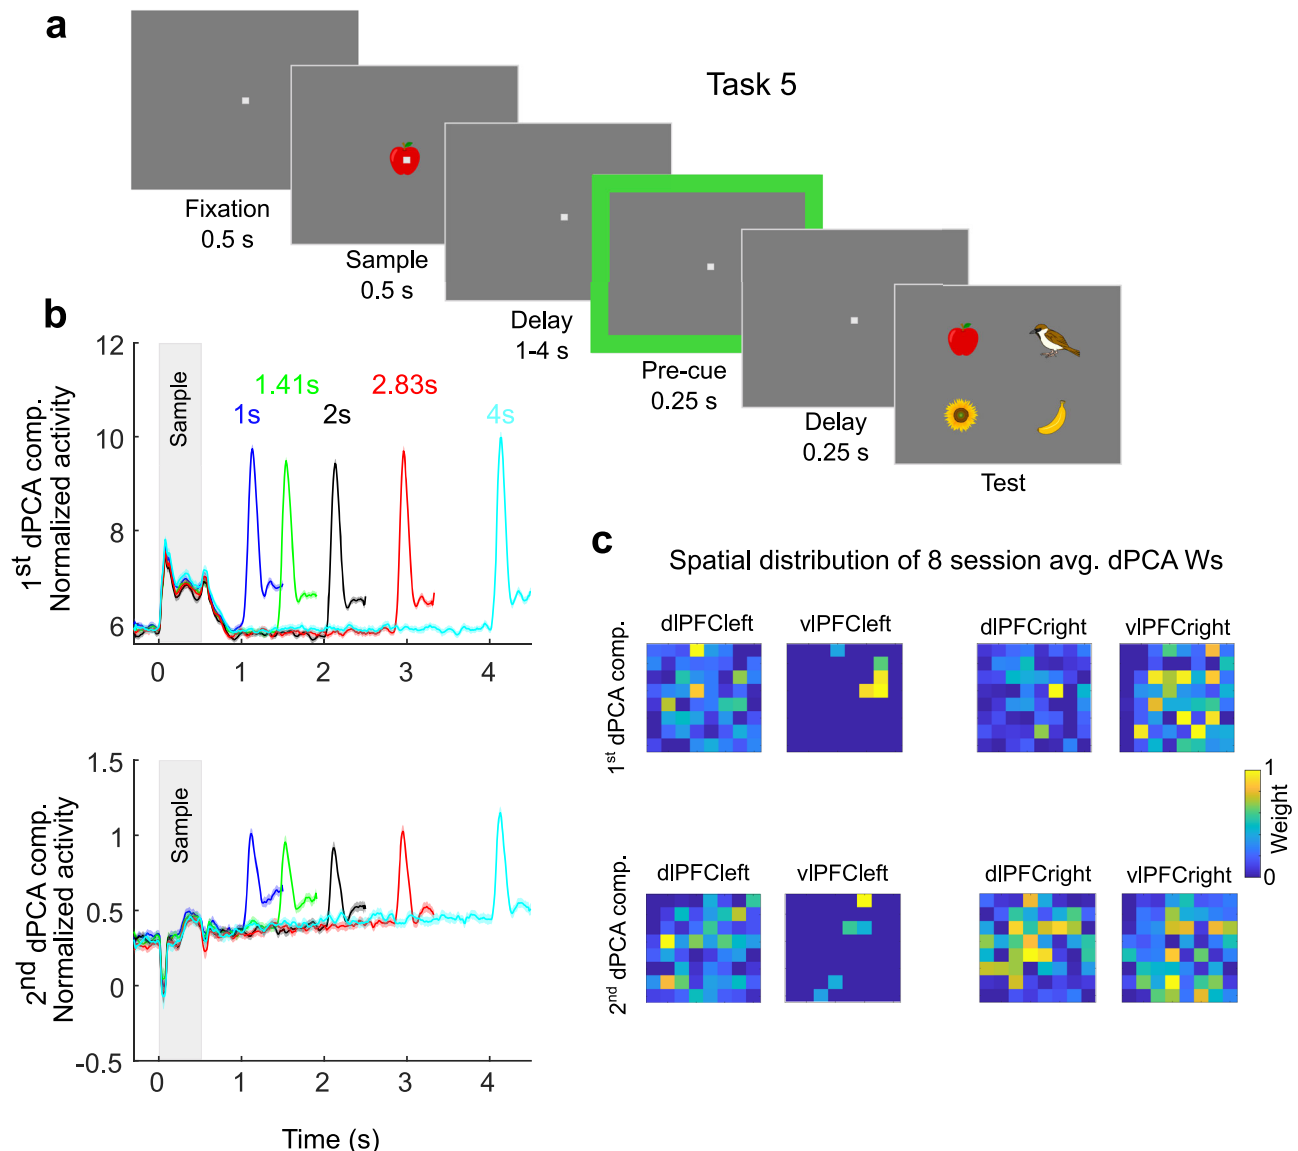


Fig. 5 | Spatial patterns of control-related activity in task with cued read-out.

a Schematic depiction of the delayed match-to-sample task with read-out cue. The delay lengths were randomly drawn from 5 possible delay lengths (1, 1.41, 2, 2.83, and 4 s). On 90% of trials, there was a test pre-cue signaling the upcoming test. Only those trials were analysed. Recording locations same as in Task 4. **b** The first (top) and second (bottom) task-general dPCA components of gamma activity ($n = 232$ electrodes). They are projected using the eight-session average dPCA weight

vectors extracted from 2 s (back) delay trials. Then data from 1 s (blue), 1.41 s (green), 2.83 s (red) and 4 s (cyan) are plotted using the same weight vectors. **c** The 4-array spatial distribution of the session-average normalized dPCA weights used in **b** for the first (top) and second (bottom) task dependent dPCA components. These were distinct from the ones obtained in Task 4 (see Fig. 4c). Source data are provided as Source Data file. Panel a was created with clip art images from © DESIGNALIKIE, Limited.

temporal patterns of gamma are activated for the first vs second item. The gamma patterns are determined by patterns in beta activity (the top-down control signal): gamma is where beta is not. When the first object is shown (say, object A), the underlying network has a unique gamma pattern corresponding to the “first item”. This injects excitation and primes neurons that are selective to A only in those gamma patches, thereby dictating the high-dimensional pattern that represents the first item. When the second object (say, B) is shown, top-down control will form a different gamma pattern corresponding to the “second item”. Neurons selective for B are then excited and primed in those patches. For the output control, where the first or second object has to be read out, the corresponding “read-out” gamma pattern is activated. Spiking of the primed neurons from that patch (network location) is elevated. Likewise, the first and second test objects have their unique low-dimensional spatial patterns of activity. The concept rests on the assumption that there is significant

redundancy in cortex such that there are neurons selective to both A and B in both patterns of patches (corresponding to the first and second object²⁴). In short, spatial computing posits that WM control stems from spatio-temporal activity patterns across network space that reflect and change with top-down task demands.

This separation of control (in this case, order) from content (objects A vs B) endows WM with the ability to generalize and with the transfer learning to new items not used when a given task was initially learned, i.e., zero-shot learning²⁵. This can be explained by different sources for these signals. Cortical beta oscillations are thought to emerge from loops between thalamus, cortex and subcortical structures²⁶. This suggests that top-down information is imposed, at least in part, from outside the local PFC cortical network itself^{25,27,28}. It would be shaped by reinforcement learning in this loop as a task is learned²⁵. The top-down excitation would selectively support stable retention of information in the cortical parts of the network it targets.

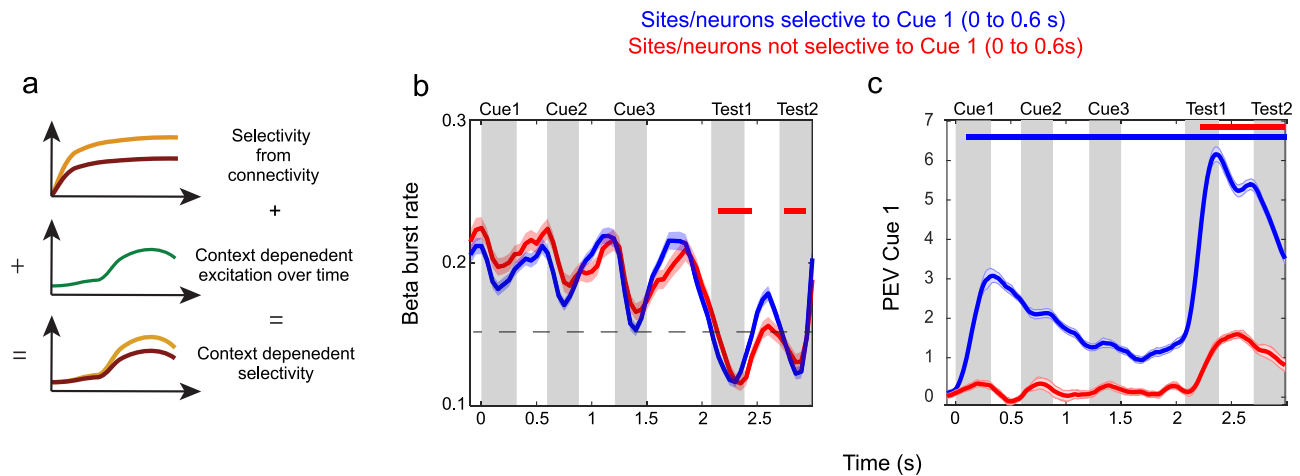


Fig. 6 | Control-related (task demand dependent) selectivity arising from time-varying excitation. **a** The conceptual illustration how in the proposed spatial computing framework mixed selectivity arises from a combination of item selectivity (orange vs brown curves, given by network connectivity) and the unique pattern of local excitation (green curve showing the control-related local excitation levels over time that are unique to each cortical location and not dependent on the item held in WM). **b** Beta bursting reflects (is anti-correlated with) the excitation in the network. Blue curves ($n = 64$) correspond to sites in which neurons are selective to items cued at Sample 1 (measured between 0 and 0.6 s), red curves ($n = 198$) describe sites with no such selective neurons. Trials with three sequential sample cues and two test cues (monkeys correctly respond at Test 2) are shown. Dotted line corresponds to the lowest beta burst rate for selective sites when not including the Test periods. Red bar

denotes times where beta burst rate in non-selective sites drops below that rate (cluster based statistic, $p < 0.05$). Grey rectangles denote cue presentations. **c** PEV in spiking accounting for the information about an item presented at Sample 1 measured in neurons on sites that are item selective (blue, $n = 69$) and non-selective (red, $n = 411$) during the first sample and delay period (0–0.6 s). Blue (selective sites) and red (non-selective sites) bars denotes times where PEV is significantly above 0 (cluster based statistic, $p < 0.05$). An exclusion criterion was adopted where data from two recording sessions were excluded due to the lack of correct trials for one of the task conditions, which would have introduced bias in the estimate of PEV information. Shaded regions indicate standard error of the mean in all panels. Source data are provided as Source Data file.

In support of this view, there is growing evidence suggesting that excitation from mediodorsal thalamus is needed to sustain working memory and attention activity in PFC^{27–29}. The selectivity of individual neurons for specific items, by contrast, may be more dependent on patterns of recurrent connectivity and inputs local to cortex^{24,30–32}.

We found evidence supporting the spatial computing hypothesis. The spatial patterns of beta and gamma reflected control operations of the task at hand. This included item order, how each item was currently being used (encoding/maintenance vs ‘reading out’ the items to make a match judgement), and which of two WM tasks the monkey was performing. Notably, the two tasks were identical except for how WM items would be read out and used at the end of the trial (recognizing a match vs choosing matches from a number of alternatives). Consistent with the spatial encoding of control signals, control-related information was found in the pattern of gamma/beta activity, shared among neurons within a few hundred micrometers. Consequently, it arose on a larger spatial scale than the item-specific activity that had more of a salt-and-pepper, finer-grained, distribution. This directly implies that the spatial dimensions of cortex are used and that oscillations can be used to selectively control information despite acting on millions of neurons simultaneously. Importantly, we observed that gamma and spiking, though highly correlated in time and space, carried different information. The gamma pattern reflected top-down control information per se. By contrast, spiking carried information about specific WM items as well as top-down information. The top-down information was ‘inherited’ by virtue of which gamma patch each neuron belonged to.

By examining recordings from chronically implanted electrode arrays, we also found that the low-dimensional gamma and beta patterns were consistent across recording sessions. Within a task, they were dynamic, changing with current task demands (meaning that different spatial parts of cortex contributed to different aspects of the task). Nonetheless, these spatio-temporal dynamics were similar across multiple recording sessions that used different WM items. We

previously reported that a ramp-up of gamma bursting was related to read-out from WM (Lundqvist et al. 2016; 2018). Here, we showed that the same parts of the network (recording array) showed gamma bursting ramp-up even though those sessions used different WM items. Spiking during these ramp-ups carry information about the specific item being read out^{11,19}. Thus, our current results suggest that the same network locations in the PFC are used for the same operation (read-out) regardless of item identity.

Spatial computing is consistent with, and sheds new light on, a variety of observations in the extant literature. Spatial computing requires representational redundancy. Information about a given WM item is represented in multiple parts of a network (for different operations). Such redundancy seems to be a hallmark of cortical activity^{33–36}. It could be supported by horizontal connections between neurons with shared stimulus preference^{24,30,37}. Spatial computing is also in line with growing interest in the dimensionality of cortical representations^{20,36,38–43}. Low-dimensional activity, i.e., shared across many neurons and across experimental conditions (such as our patterns of gamma and beta patches), has been implicated in the ability to generalize across tasks (ref.³⁹). Indeed, low-dimensional activity often reflects the structure of tasks, e.g., modulation of activity across different task periods⁴². Changes in low-dimensional activity have been shown to correlate with task learning and reflect time within a trial^{40–43}. Our findings add that these arise from spatial patterns of excitation. Spatial computing is also consistent with observations that population spiking accounting for WM items rotate into distinct subspaces depending on whether the item is currently behaviorally relevant⁴⁴. Spatial computing suggests that these rotations are driven by the spatial flow of information in networks. Moving the information would cause a rotation into a new subspace in this view.

Spatial computing can also explain classic observations that the spiking of individual neurons is highly task-dependent²⁴. Some neurons, for example, only respond to an item when it is a to-be-remembered sample or only when it is a test item used to compare

against the memorised content. Spatial computing explains this by having different sets of gamma patterns being activated for these different task contexts (see Figs. 4–6). The neurons only spike to a preferred item when its gamma patch is active. Similar results have been found in models of artificial recurrent neural networks trained on WM tasks. The artificial network units form multiple functional neuronal sub-groups with some units activating to an item when it was a sample cue, others – when it was a test item⁴⁵. Here we found similar sub-groups experimentally with the addition that they were spatially organized. This, in turn, may also provide insights into mixed selectivity^{38,46}. Neurons with mixed selectivity show context-dependent spiking that is non-linear. It cannot be predicted from their responses to the individual elements that combine to make that context. Mixed selectivity is thought to add computational horsepower, increase network storage capacity among other functional benefits^{38,46}. Here we add that the context-dependent activity of mixed selectivity neurons arises from the top-down control of low-dimensional patterns applied to networks, not just from the detailed connectivity within the network. Thus the computational benefits of mixed selectivity may be flexibly adapted from task to task.

In sum, spatial computing postulates a novel mechanism for how neural oscillations may implement selective WM control. It can also explain how PFC networks may consequently achieve flexible WM with powerful generalization capabilities. In doing so, it offers a new perspective on the functional role of low-dimensional activity that often seems to dominate cortical activity.

Methods

Data from previous studies

We analyzed data from two previous studies^{13,19}. In total the two studies included three experimental tasks (Task 1 & 2 from¹³, Task 3 from¹⁹). For details of training and data collection, please see those studies. Briefly, each task involved two Rhesus macaques that were trained until they performed well above chance. They were trained with positive reward (juice) only and maintained in accordance with the National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee for Animal Care.

For each recording, a new set of acute electrode pairs (tungsten, epoxy-coated, FHC) was lowered through a grid. Between 8 and 20 prefrontal electrodes were recorded from simultaneously on each session (34 sessions for Task 1 and 2, 30 sessions for Task 3). Task 1 and Task 2 were recorded during the same sessions in a blocked design. Only electrodes containing isolatable units were kept for further analysis.

Data recordings

We recorded data from one Rhesus macaque monkey performing two delayed-match-to sample tasks (Task 4 and Task 5). It was trained with positive reward (juice) and maintained in accordance with the National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee for Animal Care. It had 4 Utah arrays (64 channels each) chronically implanted in left vIPFC, left dIPFC, right vIPFC and right dIPFC. We recorded at a sampling rate of 30 kHz. We recorded from 5 sessions for Task 4 and 8 sessions for Task 5. For Task 4 there were 8 possible objects to be held in working memory each session, and 3/8 of the remaining objects acted as distractor at test. Sessions 2 and 3 had the same set of possible objects, and sessions 4 and 5 also shared the same set, otherwise there was no overlap across sessions (in other words, there were 3 unique sets). Over the 5 sessions the monkey performed at a high accuracy (94%), not including fixation breaks. We only analyzed the correct trials (2758, 2526, 2413, 2110, and 2306 in the 5 sessions). For Task 5 there was a pre-test cue in 90% of trials alerting the monkey to the upcoming test. We only used data from correct trials with the pre-test cue. There were 5 possible delay lengths (1, 1.41, 2,

2.83, and 4 s) to keep the timing of the test unpredictable and encourage the monkey to use the pre-test cue. We used the trials with 2 s delays to extract coefficients for the dPCA (see below) and then applied the spatial dPCA filter to the data from trials with the 4 remaining delay lengths.

Signal processing

Preprocessing Task 1 and 2: At first, all electrodes without any isolatable neurons were removed. Then, a notch filter with constant phase across a session was applied to remove 60 Hz line noise and its second harmonic. On some sessions there were high-power, broadband frequency artifacts; these sessions were discarded from further analysis.

Preprocessing Task 3: 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 (from 30 kHz).

Preprocessing Task 4 and 5: We removed channels that had bad contacts and much higher (noise) amplitudes than the rest using visual inspection. Data was downsampled to 1 kHz and 60 Hz noise removed.

For spectral analysis we applied multi-taper analysis (with a family of orthogonal tapers produced by Slepian functions^{47–49}). 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.

Burst extraction

To extract bursts of high-power events on a single trial level we utilized a previously developed method^{11,19}. 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. 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 single-trial point process (binary state: no burst vs burst within a 10 ms window) with the resolution of 10 ms and 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).

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 statistics⁵⁰, 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.

Estimation of information

The bias-corrected PEV³¹ was estimated across trials with different conditions from firing rates averaged in 50 ms bins across trials within each trial. We performed two-way ANOVA where trials had multiple groupings (i.e. stimulus or delay/task). All correct trials were used, as the groups were well balanced in each session. The bias correction was used as it avoids the problem of non-zero mean PEV for small sample sizes.

As a result, (bias-corrected) PEV allowed for the quantification of information carried by the modulation of firing rates or burst rates of individual units accounting for the stimulus, task or task period (pre-Test 1 delay vs pre-Test 2 delay in Fig. 1).

For Fig. 6 we used PEV to estimate selective and non-selective units. Here, for a unit to be classified as selective it needed to have a p -value in the ANOVA test for Stimulus 1 < 0.05 . We estimated selective units based on the 300 ms of Sample 1 presentation and the following 300 ms of delay. We did not use the full delay up until test since we wanted to compare the selectivity to that during and following the Test 1 on equal footing (which had 300 ms of test period and a 300 ms delay before Test 2). Another reason to avoid using the full delay was to demonstrate that the addition of selective units in the previously unselective population was not simply due to passage of time but timed to test onset. We used all correct trials in which Sample 1 was equal to Test 1 (to increase the statistical power we used both load 2 and load 3 trials combined since they had the same Sample 1 and Test 1 timings). In these trials the same items were presented in Test 1 and Sample 1, and the monkeys responded neither during Test 1 nor the following delay but instead awaited Test 2.

Demixed principal component analysis

To identify low-dimensional manifold for neural activity, we performed a demixed principal components analysis (dPCA)²⁰. This approach allows not only for compressing the data, similarly to PCA, but also separates the underlying components with respect to the requested task parameters by demixing the dependencies of the population activity on the task parameters. In a nutshell, demixing is achieved by minimising the reconstruction error between the projections and the neural activity averaged over trials (unlike in PCA where the reconstruction error on single trials is minimised) and over the requested task parameters. In addition, when compared to PCA the method used here benefits from greater flexibility offered by using two different linear mappings for encoding vs decoding. More technical as well as theoretical details of dPCA can be found in²⁰.

In our analyses dPCA was applied to both spiking data (firing rates obtained by convolving the spike point process with 50 ms wide Gaussian kernel) and oscillatory bursts in beta and gamma bands (burst point process convolved with 50 ms wide Gaussian kernel). To achieve demixing effect we grouped trials into task- (Task 1 vs Task 2) and stimulus- (item identity in Sample 1) dependent sets, and analyzed trials in the interval from 100 ms prior to the first sample cue (Sample 1) until the first test cue (Test 1). Apart from task and stimulus-dependent components, dPCA also produced a condition-independent component corresponding to low-dimensional time-dependent task activity.

We also used dPCA to study spatial distribution of LFP activity patterns. To this end, dPCA loadings (weights) were extracted for the condition-independent components in each session and, to avoid

overfitting to noise, a greedy search for the minimal set of weights that preserve the original component, i.e., with the mean square error below the fixed threshold, was employed. As a result, the weight vectors were made sparse by effectively removing the contribution from 42% of electrodes on average (ranging from 34 to 55%). The saliency of the contribution of particular electrodes to each component was attributed to the absolute value of the corresponding loading coefficient in the reduced weight vector.

Ethics statement

This study was performed in strict accordance with the National Institutes of Health guidelines and the policies of the Massachusetts Institute of Technology Committee for Animal Care.

Reporting summary

Further information on research design is available in the Nature Portfolio Reporting Summary linked to this article.

Data availability

The data that support the findings of this study are available from one of the authors (E.K.M.) upon reasonable request. Source data are published online with the paper. Source data are provided with this paper.

Code availability

Key functions for estimating a percentage of explained variance (PEV) and oscillatory burst analysis are available online (<https://tinyurl.com/ys7psv5u>). External tools such as Fieldtrip toolbox (Oostenveld et al. 2011) used for spectral analysis and dPCA toolbox (Kobak et al. 2016) are published online.

References

1. Baddeley, A. Working memory. *Science* **255**, 556–559 (1992).
2. D'Esposito, M. et al. The neural basis of the central executive system of working memory. *Nature* **378**, 279–281 (1995).
3. D'Esposito, M., Postle, B. R. & Rypma, B. Prefrontal cortical contributions to working memory: Evidence from event-related fMRI studies. *Exp. Brain Res.* **133**, 3–11 (2000).
4. DeNicola, A. L., Park, M. Y., Crowe, D. A., MacDonald, A. W. & Chafee, M. V. Differential roles of mediodorsal nucleus of the thalamus and prefrontal cortex in decision-making and state representation in a cognitive control task measuring deficits in schizophrenia. *J. Neurosci.* **40**, 1650–1667 (2020).
5. Goldman-Rakic, P. S. Cellular basis of working memory. *Neuron* **14**, 477–485 (1995).
6. Miller, E. K. & Cohen, J. D. An integrative theory of prefrontal cortex function. *Annu Rev. Neurosci.* **24**, 167–202 (2001).
7. Oberauer, K. Access to information in working memory: Exploring the focus of attention. *J. Exp. Psychol. Learn Mem. Cogn.* **28**, 411–421 (2002).
8. Sauseng, P., Klimesch, W., Schabus, M. & Doppelmayr, M. Frontoparietal EEG coherence in theta and upper alpha reflect central executive functions of working memory. *Int J. Psychophysiol.* **57**, 97–103 (2005).
9. Chatham, C. H., Frank, M. J. & Badre, D. Corticostriatal output gating during selection from working memory. *Neuron* **81**, 930–942 (2014).
10. Chatham, C. H. & Badre, D. Multiple gates on working memory. *Curr. Opin. Behav. Sci.* **1**, 23–31 (2015).
11. Lundqvist, M., Herman, P., Warden, M. R., Brincat, S. L. & Miller, E. K. Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nat. Commun.* **9**, 394 (2018).
12. Lewis-Peacock, J. A., Drysdale, A. T. & Postle, B. R. Neural evidence for the flexible control of mental representations. *Cereb. Cortex* **25**, 3303–3313 (2015).

13. Warden, M. R. & Miller, E. K. Task-dependent changes in short-term memory in the prefrontal cortex. *J. Neurosci.* **30**, 15801–15810 (2010).
14. Wolff, M. J., Jochim, J., Akyürek, E. G. & Stokes, M. G. Dynamic hidden states underlying working-memory-guided behavior. *Nat. Neurosci.* **20**, 864–871 (2017).
15. Yu, Q., Teng, C. & Postle, B. R. Different states of priority recruit different neural representations in visual working memory. *PLoS Biol.* **18**, e3000769 (2020).
16. van Ede, F., Niklaus, M. & Nobre, A. C. Temporal expectations guide dynamic prioritization in visual working memory through attenuated α oscillations. *J. Neurosci.* **37**, 437–445 (2017).
17. Miller, E. K., Lundqvist, M. & Bastos, A. M. Working memory 2.0. *Neuron* **100**, 463–475 (2018).
18. Lundqvist, M., Herman, P. & Lansner, A. Theta and gamma power increases and alpha/beta power decreases with memory load in an attractor network model. *J. Cogn. Neurosci.* **23**, 3008–3020 (2011).
19. Lundqvist, M. et al. Gamma and beta bursts underlie working memory. *Neuron* **90**, 152–164 (2016).
20. Kobak, D. et al. Demixed principal component analysis of neural population data. *eLife* **5**, e10989 (2016).
21. Sawaguchi, T. & Iba, M. Prefrontal cortical representation of visuospatial working memory in monkeys examined by local inactivation with muscimol. *J. Neurophysiol.* **86**, 2041–2053 (2001).
22. Hagler, D. J. Jr. & Sereno, M. I. Spatial maps in frontal and prefrontal cortex. *Neuroimage* **29**, 567–577 (2006).
23. Ray, S. & Maunsell, J. H. Different origins of gamma rhythm and high-gamma activity in macaque visual cortex. *PLoS Biol.* **9**, e1000610 (2011).
24. Goldman-Rakic, P. S. Regional and cellular fractionation of working memory. *Proc. Natl Acad. Sci. USA* **93**, 13473–13480 (1996).
25. O'Reilly, R. C. & Frank, M. J. Making working memory work: A computational model of learning in the prefrontal cortex and basal ganglia. *Neural Comput.* **18**, 283–328 (2006).
26. Ketz, N. A., Jensen, O. & O'Reilly, R. C. Thalamic pathways underlying prefrontal cortex–medial temporal lobe oscillatory interactions. *Trends Neurosci.* **38**, 3–12 (2015).
27. Schmitt, L. I. et al. Thalamic amplification of cortical connectivity sustains attentional control. *Nature* **545**, 219–223 (2017).
28. Lusk, N., Meck, W. H. & Yin, H. H. Mediodorsal thalamus contributes to the timing of instrumental actions. *J. Neurosci.* **40**, 6379–6388 (2020).
29. Bolkan, S. S. et al. Thalamic projections sustain prefrontal activity during working memory maintenance. *Nat. Neurosci.* **20**, 987 (2017).
30. Ko, H. et al. The emergence of functional microcircuits in visual cortex. *Nature* **496**, 96–100 (2013).
31. Kritzer, M. F. & Goldman-Rakic, P. S. Intrinsic circuit organization of the major layers and sublayers of the dorsolateral prefrontal cortex in the rhesus monkey. *J. Comp. Neurol.* **359**, 131–143 (1995).
32. Rao, S. G., Williams, G. V. & Goldman-Rakic, P. S. Isodirectional tuning of adjacent interneurons and pyramidal cells during working memory: evidence for microcolumnar organization in PFC. *J. Neurophysiol.* **81**, 1903–1916 (1999).
33. Pasternak, T. & Greenlee, M. W. Working memory in primate sensory systems. *Nat. Rev. Neurosci.* **6**, 97–107 (2005).
34. Siegel, M., Buschman, T. J. & Miller, E. K. Cortical information flow during flexible sensorimotor decisions. *Science* **348**, 1352–1355 (2015).
35. Dotson, N. M., Hoffman, S. J., Goodell, B. & Gray, C. M. Feature-based visual short-term memory is widely distributed and hierarchically organized. *Neuron* **99**, 215–226 (2018).
36. Stringer, C. et al. Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364**, 255 (2019).
37. Kisvárdy, Z. F., Toth, E., Rausch, M. & Eysel, U. T. Orientation-specific relationship between populations of excitatory and inhibitory lateral connections in the visual cortex of the cat. *Cereb. Cortex* **7**, 605–618 (1997).
38. Rigotti, M. et al. The importance of mixed selectivity in complex cognitive tasks. *Nature* **497**, 585–590 (2013).
39. Badre, D., Bhandari, A., Keglovits, H. & Kikumoto, A. The dimensionality of neural representations for control. *Curr. Opin. Behav. Sci.* **38**, 20–28 (2020).
40. Cueva, C. J. et al. Low dimensional dynamics for working memory and time encoding. *Proc. Natl Acad. Sci. USA* **117**, 23021–23032 (2019).
41. Wolff, M. J., Jochim, J., Akyürek, E. G., Buschman, T. J. & Stokes, M. G. Drifting codes within a stable coding scheme for working memory. *PLoS Biol.* **18**, e3000625 (2020).
42. Tang, H., Qi, X. L., Riley, M. R. & Constantinidis, C. Working memory capacity is enhanced by distributed prefrontal activation and invariant temporal dynamics. *Proc. Natl Acad. Sci. USA* **116**, 7095–7100 (2019).
43. MacDowell, C. J. & Buschman, T. J. Low-dimensional spatio-temporal dynamics underlie cortex-wide neural activity. *Curr. Biol.* **30**, 2665–80.e8 (2020).
44. Panichello, M. F. & Buschman, T. J. Shared mechanisms underlie the control of working memory and attention. *Nature* **592**, 601–605 (2021).
45. Dubreuil, A., Valente, A., Beiran, M., Mastrogioseppe, F. & Ostojic, S. The role of population structure in computations through neural dynamics. *Nat. Neurosci.* **25**, 783–794 (2020).
46. Fusi, S., Miller, E. K. & Rigotti, M. Why neurons mix: High dimensionality for higher cognition. *Curr. Opin. Neurobiol.* **37**, 66–74 (2016).
47. Slepian, D. Prolate spheroidal wave functions, Fourier analysis, and uncertainty—V: The discrete case. *Bell Syst. Tech. J.* **57**, 1371–1430 (1978).
48. Thomson, D. J. Spectrum estimation and harmonic analysis. *Proc. IEEE* **70**, 1055–1096 (1982).
49. Jarvis, M. R. & Mitra, P. P. Sampling properties of the spectrum and coherency of sequences of action potentials. *Neural Comput.* **13**, 717–749 (2001).
50. Maris, E. & Oostenveld, R. Nonparametric statistical testing of EEG- and MEG-data. *J. Neurosci. Methods* **164**, 177–190 (2007).
51. Olejnik, S. & Algina, J. Generalized eta and omega squared statistics: measures of effect size for some common research designs. *Psychol. Methods* **8**, 434 (2003).

Acknowledgements

We would like to thank the ERC starting grant 949131, the Swedish Research Council Starting Grant 2018-04197, the Swedish Research Council grants: 2018-05360 and 2016-05871, the 2017 Young Investigator Grant from the Brain & Behavior Research Foundation, the National Institutes of Mental Health Grant R37MH087027 and 5K99MH116100-02, Office of Naval Research N000142212453, The JPB Foundation, The Picower Institute for Learning and Memory, and Digital Futures and Swedish e-Science Research Center for their support.

Author contributions

Formal analysis: P.H. and M.L.; Investigation: M.L., S.B., T.B., J.R., and M.W.; Data curation: S.B. and P.H.; Writing: M.L., P.H., and E.K.M.; Validation: P.H.; Conceptualization: M.L., P.H., and E.K.M.; Resources: E.K.M.; Visualization: M.L.; Funding acquisition: M.L., E.K.M., and P.H.

Funding

Open access funding provided by Royal Institute of Technology.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41467-023-36555-4>.

Correspondence and requests for materials should be addressed to Mikael Lundqvist or Pawel Herman.

Peer review information *Nature Communications* thanks the anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at <http://www.nature.com/reprints>

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this license, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2023