Oscillatory Control of Cortical Space as a Computational Dimension

Zhen Chen[1], Scott L. Brincat[1], Mikael Lundqvist[1,2], Roman F. Loonis[1,3], Melissa R. Warden[1,4], Earl K. Miller[1]*

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

[2] Department of Clinical Neuroscience, Karolinska Institutet, Stockholm, Sweden

[3] Takeda Pharmaceuticals America, Cambridge, MA 02142, USA

[4] Department of Translational Neurosciences, University of Arizona College of Medicine - Phoenix, Phoenix, AZ, USA and Graduate Interdisciplinary Program in Neuroscience, University of Arizona, Tucson, AZ, USA

* Corresponding author

Email: ekmiller@mit.edu

Abstract

Flexible cognition depends on the ability to represent and apply context, allowing the brain to interpret sensory input and guide behavior in a context-dependent manner. Recent work has proposed *Spatial Computing* as a mechanism for this flexibility, suggesting that contextual signals organize information processing through spatial patterns of oscillatory activity across the cortical surface. These patterns act as “inhibitory stencils” that constrain where information (the “content” of cognition) can be expressed in spiking activity. Here, we provide a comprehensive empirical test of Spatial Computing Theory using multi-electrode recordings from the lateral prefrontal cortex in non-human primates performing a range of cognitive tasks (object working memory, sequence working memory, categorization). We found that alpha/beta oscillations encoded contextual information, reorganized their spatial patterns with context and task demands, and spatially gated the expression of content-related spiking activity. Furthermore, we found that alpha/beta oscillations reflected misattributions of task context and correlated with subjects’ trial-by-trial decisions. These findings validate core predictions of Spatial Computing by showing that oscillatory dynamics not only gate information in time but also shape where in the cortex cognitive content is represented. This framework offers a unifying principle for understanding how the brain flexibly coordinates cognition through structured population dynamics.

Introduction

A hallmark of cognition is its flexibility, the ability to adjust thoughts and actions according to current goals, rules, and contexts. This flexibility depends on more than just storing information. It requires the brain to control the expression of relevant information, when it should be accessed, and how it should be used. Central to this capacity is the representation and use of *context*. It shapes how sensory input (i.e., *content*) is interpreted and how behavior is guided.

Recent work has proposed a novel principle—*Spatial Computing*—as a potential mechanism for flexible context-dependent control (*1*, *2*). The core idea is that the brain achieves flexible control by using oscillatory dynamics to control neural activity across the cortical surface. In essence, cortical space is utilized as a computational dimension. In this framework, signals reflecting cognitive context are carried by oscillatory activity, particularly alpha/beta-band oscillations (10–30 Hz). By exerting inhibition, these low-dimensional mesoscale patterns modulate local neuronal excitability and determine which high-dimensional populations of neurons on the microscale can be activated in a given context. This shapes where in the cortex content-related spiking activity can occur. Rather than relying solely on fixed anatomical pathways, the brain dynamically organizes information flow through spatially structured oscillatory patterns. In this way, Spatial Computing offers a mechanism for how information and content representations are processed and organized in a context-dependent manner.

This framework provides a mechanistic explanation for flexible control of cognition. For example, it helps explain how low-dimensional subspaces encoding different task variables can emerge, allowing the brain to reduce interference and flexibly switch between representations (*3*, *4*). It also accounts for how nonlinear mixed selectivity of neurons arises from the interaction of cognitive control and stimulus features (*5*–*8*). By dynamically modulating local neuronal excitability, alpha/beta oscillatory patterns allow the same neural populations to support different cognitive operations depending on context.

Despite these advances, several critical questions about Spatial Computing remain unanswered. First, the theory proposes that different neural signals serve distinct roles—oscillatory patterns reflect cognitive *control* and spiking activity represents the *contents* of cognition. However, little work has been done to directly compare the information carried by each signal type at the *population* level. Second, although Spatial Computing proposes that contextual signals are reflected in alpha/beta oscillations, it is still unknown whether the spatial patterns of alpha/beta power reorganize with context. It is also unknown whether any such reorganization between contexts scales with cognitive demands, i.e. stronger cognitive demands result in larger reorganization between contexts. Third, while the theory posits that alpha/beta regulates where content is expressed, no direct test has shown that spiking-based content representation is spatially gated by alpha/beta patterns across the cortical space. Finally, although Spatial Computing is hypothesized to support flexible cognition and behavior, there is little direct evidence that alpha/beta or gamma (30–150 Hz) oscillations predict trial-by-trial decisions or reflect behavioral errors.

Together, these open questions point to the need for a more comprehensive empirical test of Spatial Computing Theory.

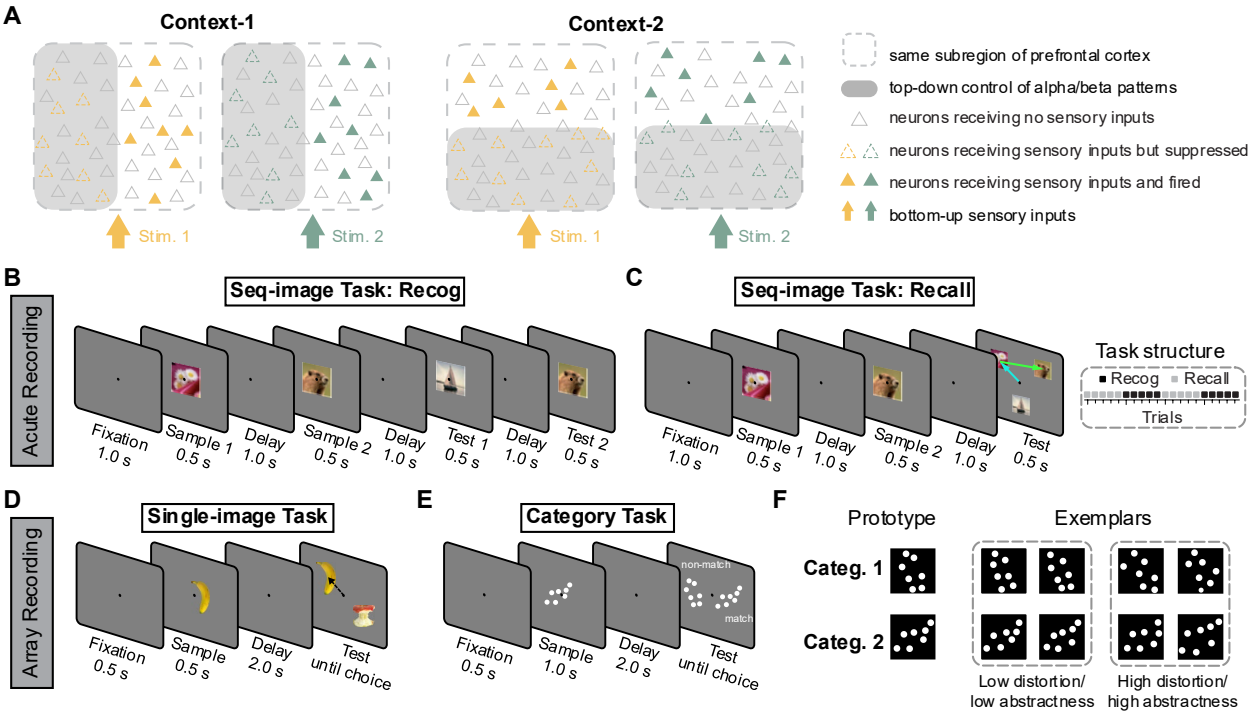

Figure 1. Schematic overview of Spatial Computing and task designs.
(A). Conceptual framework of Spatial Computing Theory. Alpha/beta oscillations convey cognitive control signals reflecting context, and their spatial pattern gates where information ("content") can be expressed by spiking. Shaded grey areas indicate stronger alpha/beta power where spiking is suppressed while unshaded areas indicate lower alpha/beta power where spiking is expressed.
(B). "Recog" version of Seq-image Task showing a sequence of sample images followed by a delay. During the test phase, NHPs judged whether a probe sequence matched the sample sequence.
(C). "Recall" version of Seq-image Task where, during the test phase, NHPs made saccades to the remembered objects in their original order. The two versions alternated in trial blocks (right).
(D). Single-image task where a single image was shown and, after a delay, NHPs made a saccade to the remembered object.
(E). Category-match task where NHPs categorized exemplar dot patterns into two learned categories.
(F). The exemplar dot patterns were generated by distorting two prototype patterns (left). Distortion levels varied the similarity between exemplars and prototypes, controlling the degree of abstractness (low vs. high).

We analyzed multi-electrode recordings from the lateral prefrontal cortex as non-human primates (NHPs) performed various cognitive tasks requiring working memory (WM), sequencing, and/or categorization. We formalized five core predictions of Spatial Computing and tested them across different cognitive demands. Our results show that alpha/beta patterns encode contextual information, flexibly reorganize with task context and cognitive demand, spatially govern where content is expressed in spiking activity, and closely track behavioral performance. These results provide empirical support for Spatial Computing Theory. They highlight its potential as a unifying framework for

understanding how the brain flexibly coordinates cognition through structured population dynamics.

Results

Predictions of Spatial Computing Theory and Tasks

We tested five basic predictions of *Spatial Computing Theory*: 1. Alpha/beta oscillations primarily convey information reflecting cognitive control and context. 2. Spiking activity conveys information about both context and content (e.g., stimulus identity). 3. The spatial patterns of alpha/beta oscillations change with context and cognitive demands. 4. Alpha/beta patterns spatially organize the expression of content information in spiking activity. 5. Alpha/beta oscillations reflect NHPs' behavioral performance and trial-by-trial choices. As downstream effectors of cognitive control and structured by alpha/beta, spiking and associated gamma are expected to correlate with behavior as well. **Figure 1A** illustrates this framework, where alpha/beta patterns carry cognitive control signals and are dynamically shaped by the current context. These alpha/beta patterns act as "inhibitory stencils" which constrain where higher levels of spiking occur. This allows the stimulus-driven content information to be expressed only in regions where alpha/beta is weak. In essence, cortical space becomes a computational dimension. Alpha/beta patterns sculpt the spatial landscape of neural activity, thereby supporting context-dependent information processing.

To test these predictions, we analyzed multi-electrode neurophysiological recordings (local field potentials (LFPs) and single/multi-unit spiking activity) from the lateral prefrontal cortex (PFC) of NHPs (rhesus macaques) performing three cognitive tasks (**Figure S1**). In one of the tasks, the Seq-image Task (*9*), the NHPs performed a sequential working memory (WM) task involving two sample items (**Figure 1B**). Two versions of the task were used. Both required working memory for two sample items and their order of appearance. They only differed in the task rule that determined how NHPs reported these memories. In the 'Recog' version, the NHPs indicated whether a subsequent sequence of test items was a match to the sample sequence (**Figure 1B1**). In the 'Recall' version, NHPs were offered three test items and asked to saccade to the two seen earlier in the correct order (**Figure 1B2**). The NHPs alternated between the two versions across blocks of trials. Thus, NHPs needed to maintain two types of contextual information: the order in which the two samples were presented on each trial, and the current task rule that needed to be applied to the test items.

The two other tasks involved working memory and categorization of individual images. The Single-image task (*10*, *11*) was a standard delayed match-to-sample working memory task (**Figure 1D**). In the Category Task (*12*, *13*), NHPs had to categorize dot-pattern stimuli into two learned categories, each defined by random distortions around a prototype (**Figure 1E1**). The level of distortion ("abstractness") was varied by manipulating the similarity between the exemplars and their corresponding prototype (**Figure 1E2**). Thus, in this task, the individual dot exemplars were the "content" of cognitive processing. Their inferred category was the context that determined which behavioral response was appropriate for each exemplar. Additionally, manipulating the level of exemplar distortion across trials allowed us to test different degrees of cognitive demands.

These three tasks allowed us to probe how cognitive content (e.g., identity of sample items and exemplar dot patterns) and context (e.g., task rule, order, and category) is encoded by spiking and LFP dynamics. By examining the spatial distribution of these signals, we could directly test whether alpha/beta patterns form a “stencil” that organizes content-related activity expressed by spiking. A complete description of the experimental design for each task can be found in the **Methods** and in the original publications that first reported these data (*1, 9–14*).

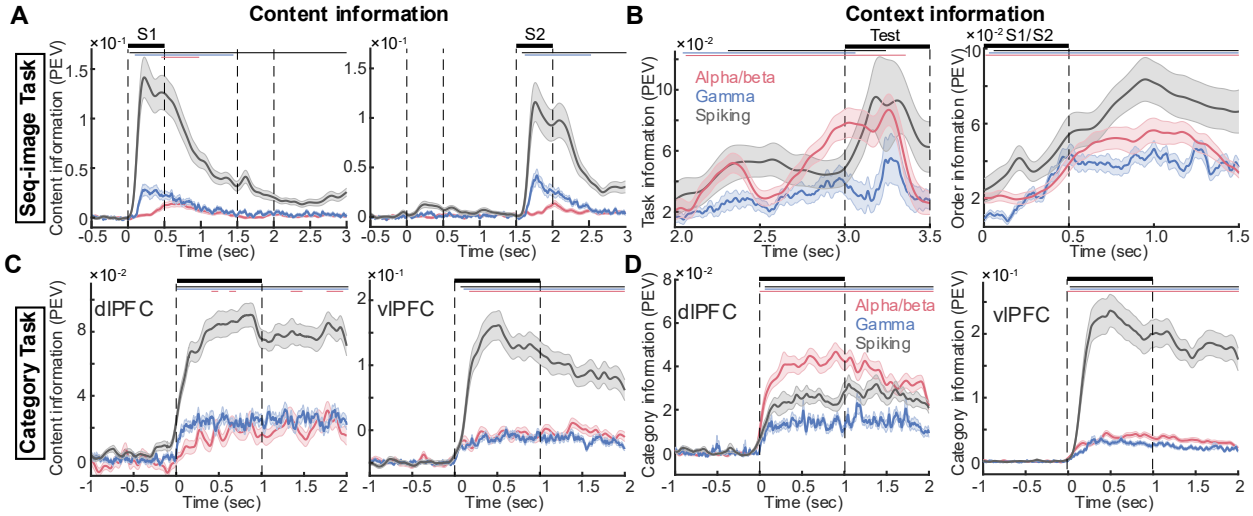

Figure 2. Alpha/beta power primarily carried context information, while spiking activity conveyed mixed content and context signals.

(A). Content information (quantified by PEV) conveyed by each signal type in the Seq-image Task, for sample 1 (left) and sample 2 (right). Vertical dashed lines indicated the sample presentation periods (S1 and S2). Horizontal thin lines on top reflect times of significant PEV values for the corresponding signal type using the same color code (p<0.001, one-sample t-test). Error bar: mean ± SE, n = 48 (2 task versions × 24 sessions).

(B). Context information conveyed by each signal type in the Seq-image Task. Left: PEV for task version (Recog vs. Recall) during the second delay and test period. Right: PEV for sequence order (first vs. second order) during the sample presentation followed by delay. Error bar: mean ± SE, n = 24 sessions. Significance: p<0.001, one-sample t-test.

(C). Content information for each signal type in the Category Task. Left: PEV for dlPFC. Right: PEV for vlPFC. Error bar: mean ± SE, n = 51 sessions. Significance: p<0.001, one-sample t-test.

(D). Category context information for each signal type in the Category Task. Left: PEV for dlPFC. Right: PEV for vlPFC. Error bar: mean ± SE, n = 51 sessions. Significance: p<0.001, one-sample t-test.

See Fig. S2 for corresponding paired comparisons between signal types.

Predictions 1 and 2: Alpha/beta Reflects Context, Spiking Reflects Context and Content

To test these predictions, we quantified the amount of information reflecting content and context carried by each signal type. Rather than analyzing single electrodes or individual units, we developed a method to compute the percentage of explained variance (PEV) by each factor across the entire population of units or LFPs (see **Methods**). For each factor, we computed an axis in population activity (spike rates or LFP band power) space that

optimally separated contrasted conditions. We then projected population activity onto this axis and quantified its PEV. This approach allowed us to capture the information encoded by population-level LFP power and spiking.

We began with the content information in the Seq-image Task—the identity of the two sample items to be remembered (**Figure 1B**). In line with previous studies (*15*), we found that spiking activity carried the most content information. Information increased sharply during each sample presentation (**Figure 2A**). There was some content information in gamma, but it was much lower than in spiking activity (**Figure S2**). Content information was minimal in alpha/beta power, in line with Prediction 1.

We then assessed context encoding by measuring information conveyed about the *task rule* being performed (i.e., the task version, Recog vs. Recall) during the delay period preceding test onset. Throughout the delay, all signals carried task information but alpha/beta and spiking carried more than gamma (**Figure 2B**, left; see **Figure S2** for paired comparison). Note the ramping up of alpha/beta task information. It peaked near the end of the delay and plateaued through the presentation of the test items when the task rules (i.e., context) needed to be applied. In addition to task rules, we also examined the sample sequence order (i.e., first vs. second sample item), which is another type of contextual factor in this paradigm. We computed information reflecting the sequence presentation order (i.e., sample/delay 1 vs. sample/delay2), while holding constant the sample image shown (see **Methods**). Similar to task information, all signals carried sequence order information, with alpha/beta and spiking carrying more than gamma (**Figure 2B**, right; see **Figure S2** for paired comparison). These results show that, in the Seq-image Task, cognitive context is reflected mainly in alpha/beta power and spiking, while WM content is reflected primarily in spiking activity, consistent with Predictions 1 and 2.

We then turned to the Category Task. The content of this task was the individual category exemplars, while the cognitive context was their inferred category membership (which determined the appropriate behavior for each exemplar). For content information, spiking activity carried significantly more information than both alpha/beta and gamma oscillations (**Figure 2C**, see **Figure S2** for paired comparison). For this dataset, we have extensive sampling of both the dorsolateral (dlPFC) and ventrolateral (vlPFC) prefrontal subregions. This pattern was seen in both dlPFC and vlPFC. By contrast, for context information, we observed a clear distinction between the two regions, as we discuss next.

Previous work has shown that the dlPFC is more associated with top-down processing, while vlPFC is relatively more driven by bottom-up sensory stimulation (*13*, *16*–*18*). Consistent with this, alpha/beta oscillations in dlPFC were the primary carriers of context information (**Figure 2D**, left; see **Figure S2** for paired comparison). By contrast, in vlPFC, spiking activity dominated (**Figure 2D**, right; see **Figure S2** for paired comparison). Importantly, these differences between dlPFC and vlPFC were not due to differences in overall spiking rate or recording quality between these two regions. Although vlPFC showed a brief (~100 ms) elevation in mean spiking rate relative to dlPFC (**Figure S3**), their spiking rates were nearly identical through the rest of the trial. This indicates the

maintained elevated information about context in vlPFC spiking was not due to overall higher spikes rates. Moreover, dlPFC and vlPFC carried comparable amounts of information reflecting the behavioral response (choice of left vs. right test exemplar; **Figure S3**). This indicates that the observed differences between dlPFC and vlPFC arose from their functional distinctions rather than differences in data quality. This set of results shows that alpha/beta mainly reflects context, while spiking activity reflects both context and content. This is consistent with Predictions 1 and 2.

For a further test of context, we examined the level of abstractness of the categories. The Category Task manipulated the level of abstractness by different degrees of distortions of the prototypes. For low distortions, the exemplars from the same category were more similar in appearance to one another and to their prototypes. Thus, they could be categorized by their bottom-up appearance, and required less abstraction. By contrast, at high distortion levels, exemplars from different categories could appear similar to each other. They thus had to be categorized by a category rule learned by the NHPs (*13*), i.e, by the cognitive context.

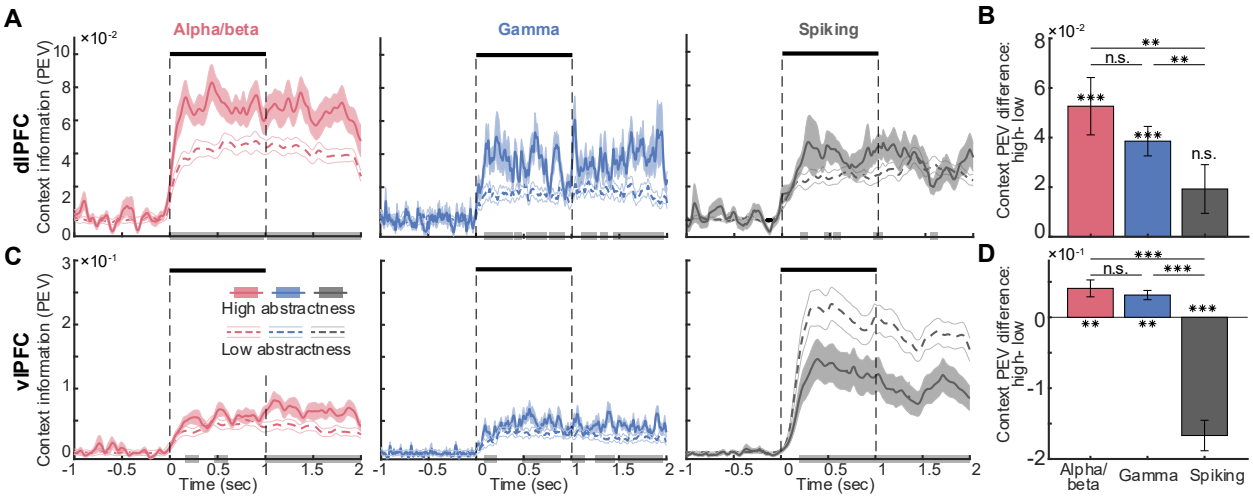

Figure 3. Category context encoding varies with abstractness level.

(A). Category context information (PEV) by alpha/beta (left), gamma (middle), and spiking (right) in dlPFC, shown separately for high- (solid lines) and low-abstractness (dashed lines) conditions. Horizontal grey bars on the x-axis represented significant PEV differences between high- and low-abstractness (p<0.05, two-sample t-test). Error bar: mean ± SE, n = 51 sessions.

*(B). Cumulative difference in category PEV (high- minus low-abstractness) across time in dlPFC, shown for each signal type. *** p<0.001, **: p<0.01, n.s.: not significant, one-sample t-test for each signal type and Wilcoxon rank-sum test across signals. Error bar: mean ± SE, n = 51 sessions.*

(C). Same as (A) but for vlPFC.

(D). Same as (B) but for vlPFC.

This revealed context information in alpha/beta was stronger for the high-abstractness exemplars (especially in dlPFC). By contrast, context information in spiking was stronger for the low-abstractness exemplars (in vlPFC). In dlPFC, alpha/beta (**Figure 3A,B**, left) and gamma (**Figure 3A,B**, center) power encoded significantly more context information under high-abstractness than under low-abstractness. These effects were greater than

the modulation of dlPFC spiking activity by abstractness, which was itself not significant (**Figure 3A,B**, right). In vlPFC, by contrast, spiking activity carried more context information under low-abstractness conditions than under high abstractness (**Figure 3C,D**, right), while alpha/beta and gamma showed the opposite direction of modulation by abstractness level (**Figure 3C,D**, left and center). Thus, when exemplars are heavily distorted (high-abstractness), dlPFC alpha/beta plays a dominant role in representing the relevant context through top-down mechanisms. Conversely, when exemplars are easily recognized (low-abstractness), vlPFC spiking can represent the category more directly through bottom-up mechanisms. This is consistent with Predictions 1 and 2.

Overall, these findings support Predictions 1 and 2 of Spatial Computing. They show that alpha/beta (especially in dlPFC) strongly reflects context, while spiking reflects both content and context. Context information carried by alpha/beta (in dlPFC) is greater when tasks require greater cognitive demands. The distinction between dlPFC and vlPFC aligns with their respective bias toward top-down vs. bottom-up processing.

Prediction 3: Alpha/Beta Spatial Patterns Change with Context and Cognitive Demand

Because the Category Task used multi-electrode Utah arrays (see **Methods**), it also allowed us to test the prediction that the *spatial patterning* of alpha/beta changes with context and cognitive demands. We next assessed how distinct alpha/beta patterns were for different categories (i.e., context) and how this category differentiation was modulated by the level of abstractness of the categories (i.e., cognitive demands). To quantify this, we measured spatial correlations of alpha/beta patterns (time-averaged during sample presentation; see **Methods**) between different categories (between-category correlation). As a reference for comparison, we also measured correlations between different exemplars within the *same* category (within-category correlation; see **Methods**). If the spatial patterns of alpha/beta are truly context-sensitive, the between-category correlation should be lower than the within-category correlation.

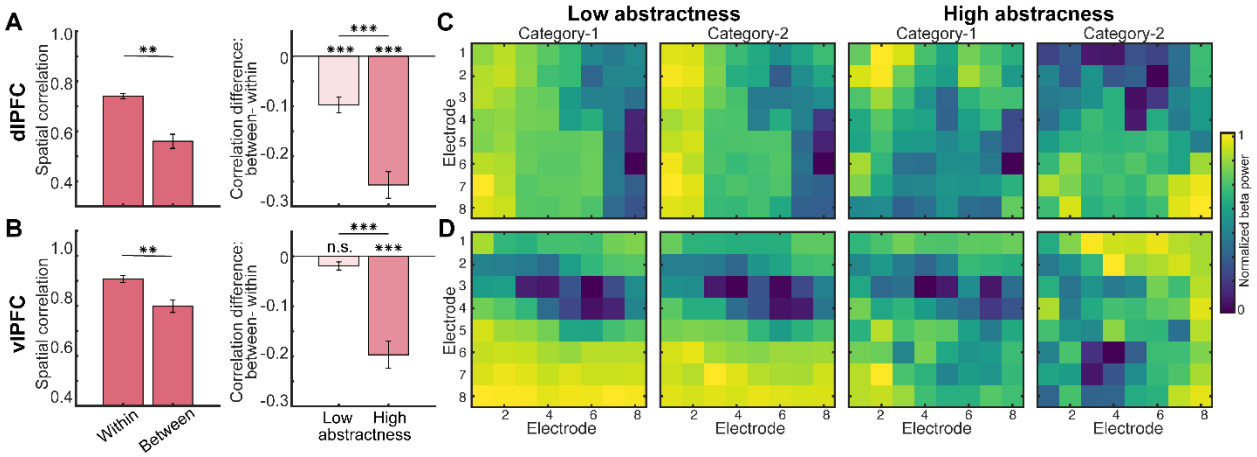

Figure 4. Spatial patterns of alpha/beta power change with category context and abstractness level.

(A). Left: overall spatial correlation of trial-averaged alpha/beta power during sample presentation, shown separately for within-category and between-category correlation in dlPFC. Right: spatial

*correlation differences (between- minus within-category) for high- and low-abstractness conditions. *** p<0.001, **: p<0.01, n.s.: not significant, two-sample t-test. Error bar: mean ± SE, n = 51 sessions.*
(B). Same as (A) but for vlPFC.
(C). Example spatial patterns of alpha/beta power averaged across trials, shown for two categories under low- and high-abstractness in one session from dlPFC.
(D). Same as (C) but for vlPFC.

In both dlPFC and vlPFC, the overall between-category correlations were significantly lower than within-category correlations (**Figure 4A,B**, left). This indicates that alpha/beta spatial patterns differentiate category membership and thus reflect context. To summarize these results, we quantified the difference in correlation between these conditions (between-category – within-category). We measured these correlation differences separately for low- and high-abstractness category exemplars (**Figure 4A,B**, right). For both dlPFC and vlPFC, this difference was significantly greater for high- than low-abstractness. This demonstrates that alpha/beta spatial patterns are more distinct when cognitive demands are greater. For dlPFC, the differences were individually significant for both low and high abstractness (**Figure 4A**, right). By contrast, vlPFC showed no significant difference under low-abstractness, but a significant separation emerged only under high-abstractness (**Figure 4B**, right). These findings suggest that when category membership is visually straightforward (low-abstractness), alpha/beta patterns in vlPFC do not distinguish between categories, reflecting a more bottom-up drive of categorization. However, when stimuli are highly distorted and thus require more top-down processing to categorize them, spatial alpha/beta patterns, particularly in dlPFC, become more distinct for each category. Example spatial maps in **Figure 4C** and **4D** illustrate this increased category differentiation under high-abstractness in dlPFC and vlPFC respectively. These results are consistent with Prediction 3.

Prediction 4: Alpha/Beta Patterns Spatially Regulate Content Information

Alpha/beta patterns are proposed to spatially regulate the expression of content information carried by spiking. **Figure 5A** shows a schematic of this hypothesis: Areas of high alpha/beta power exhibit lower content information carried by spiking, whereas areas with low alpha/beta power show higher spiking content information.

To assess this prediction, we first analyzed data from the Single-image task (**Figure 1D**), where multi-electrode Utah arrays were implanted in both dlPFC and vlPFC. We grouped LFPs/spiking by their positions in the electrode array. For each electrode, we computed the time-averaged alpha/beta power during sample presentation. We then split electrodes into two groups—those with higher or lower alpha/beta power across electrodes. Finally, we computed PEV reflecting content (i.e. sample object identity) separately in spiking and gamma for each single electrode in order to construct spatial maps of content information (PEV) (see **Methods**).

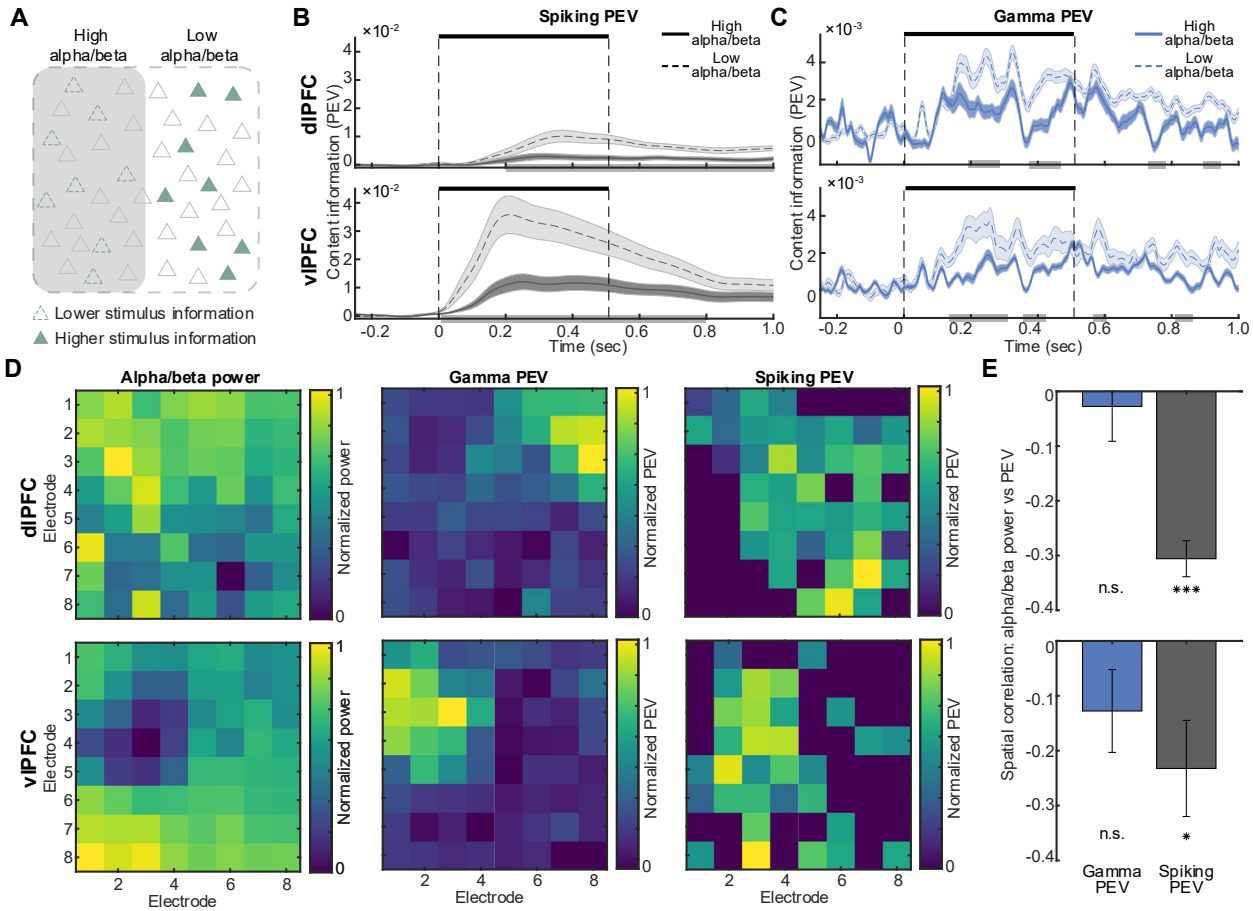

Figure 5. Alpha/beta patterns spatially regulate the distribution of content information in cortical space.

(A). Schematic illustration of spatial regulation: regions with higher alpha/beta power suppress content encoding and information in spiking.

(B). Content information (PEV for stimulus identity) carried by spiking activity at low- and high-alpha/beta electrodes in dlPFC (top) and vlPFC (bottom). Horizontal grey bars on the x-axis represented significant PEV differences between low- and high-alpha/beta power ($p<0.05$, two-sample t-test). Error bar: mean ± SE. Sample sizes (low/high alpha-beta): dlPFC spiking (n = 650/427), vlPFC spiking (n = 198/262).

(C). Same as (B) but for content PEV carried by gamma at low- and high-alpha/beta electrodes in dlPFC (top) and vlPFC (bottom). Error bar: mean ± SE. Sample sizes (low/high alpha-beta): dlPFC gamma (n = 672/730), vlPFC gamma (n = 365/909).

(D). Example spatial patterns from a single session in dlPFC (top) and vlPFC (bottom) showing time-averaged alpha/beta power, content PEV by gamma, and content PEV by spiking during the sample presentation period. Each pattern was normalized individually.

*(E). Correlation between spatial patterns of alpha/beta power and content PEV of gamma and spiking in dlPFC (top) and vlPFC (bottom). *** $p<0.001$, *: $p<0.05$, n.s.: not significant, one-sample t-test. Error bar: mean ± SE. Sample sizes: dlPFC (gamma: n = 22, spiking: n = 15), vlPFC (gamma: n = 20, spiking: n = 18).*

As predicted, spiking activity on electrodes with lower alpha/beta power exhibited significantly greater content PEV than spiking activity on electrodes with higher power (**Figure 5B**, left). This differentiation was significant but weaker in gamma (**Figure 5C**).

To visualize these spatial relationships, **Figure 5D** shows example maps of the alpha/beta power, content PEV in gamma, and content PEV in spiking from a representative session. It shows anti-correlation between the spatial patterns of alpha/beta power and of content PEV in spiking. Across all sessions, we found a significant negative spatial correlation between alpha/beta power and content PEV in spiking (**Figure 5E**). The same negative correlation was also found in the Category task, in which multi-electrode Utah arrays were used **(Figure S4**). Gamma power showed a similar, but non-significant, trend toward negative correlation with alpha/beta (**Figure 5E**). Together, these results indicate that content information is expressed strongly where alpha/beta is weak, but is weak where alpha/beta is strong. This suggests alpha/beta regulates the spatial distribution of content information, supporting Prediction 4.

Prediction 5: Alpha/Beta Correlates with Behavior

If alpha/beta truly encodes context in a way that influences the NHPs' decisions, we would expect these signals also to differ between correct and incorrect trials and to correlate with trial-by-trial decisions.

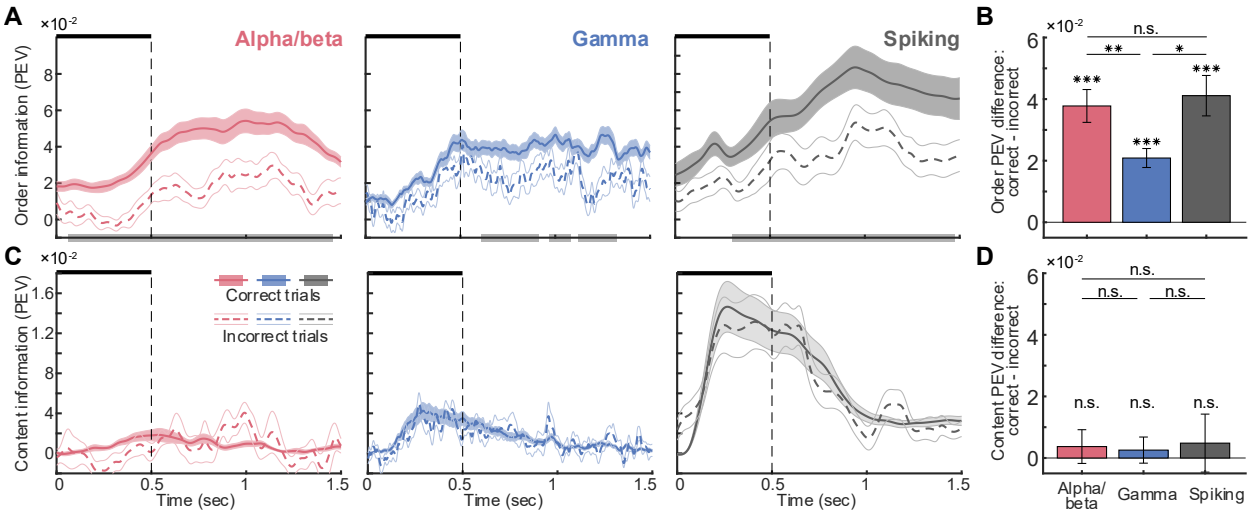

Figure 6. Alpha/beta and spiking reflect correct versus incorrect contextual decisions.

(A). Contextual information (PEV) about sample order in Seq-Image task for correct trials (solid lines) and order-error trials (dashed lines), where NHPs chose the correct samples but in reversed order. Results are shown separately for alpha/beta (left), gamma (middle), and spiking (right). Horizontal grey bars on the x-axis indicate significant PEV differences between correct and error trials (p<0.05, two-sample t-test). Error bar: mean ± SE, n = 24 sessions.

*(B). Cumulative difference in order information (correct minus incorrect) across time for each signal type. *** p<0.001, **: p<0.01, *: p<0.05, n.s.: not significant, one-sample t-test for each signal type and Wilcoxon rank-sum test for across-signal comparisons. Error bar: mean ± SE, n = 24 sessions.*

(C). Same as in (A) but for content information.

(D). Same as in (B) but for content information.

We first tested this in the Recall version of Seq-image Task, where NHPs had to report the two sample items in the correct order. This design allowed us to isolate contextual

errors. For many of the incorrect trials, NHPs chose the correct samples but in the wrong order, indicating an error of reversed context. Thus, we compared the neural information reflecting sample order between correct and order-error trials. We found that all signals carried significantly more order information in correct trials than in order-error trials (**Figure 6A**). This behavioral modulation was greatest for alpha/beta and spiking (**Figure 6B**). By contrast, information about content (i.e., sample identity) did not differ between correct and order-error trials for all signals (**Figure 6C** and **6D**). Similar results were also found for the Category Task by comparing correct vs. incorrect categorization trials (**Figure S5** and **S6**). This suggests that alpha/beta and spiking were crucial for maintaining the correct sequence order and category that influences behavioral decisions.

We next turned to the Category Task to determine whether these signals correlate with NHP's trial-by-trial decisions of category membership. We adapted a trial-by-trial analysis similar to Cohen & Maunsell, 2010(*19*). Specifically, we projected each trial's neural population data (time-averaged over sample presentation) onto an axis in population activity space that optimally separates the two learned categories on correct trials (see **Methods**). We then normalized projections onto this "category axis" so that a value of –1 corresponded to the average for Category-1 correct decisions and +1 to the average for Category-2 correct decisions. This provided a single-trial metric of population category coding strength. We then projected onto this axis the activity from the incorrect trials, which were not used to define the axis. If alpha/beta patterns are indeed relevant for categorization, we would expect the projections of the incorrect trials to be "shifted" toward the opposite, incorrect, category.

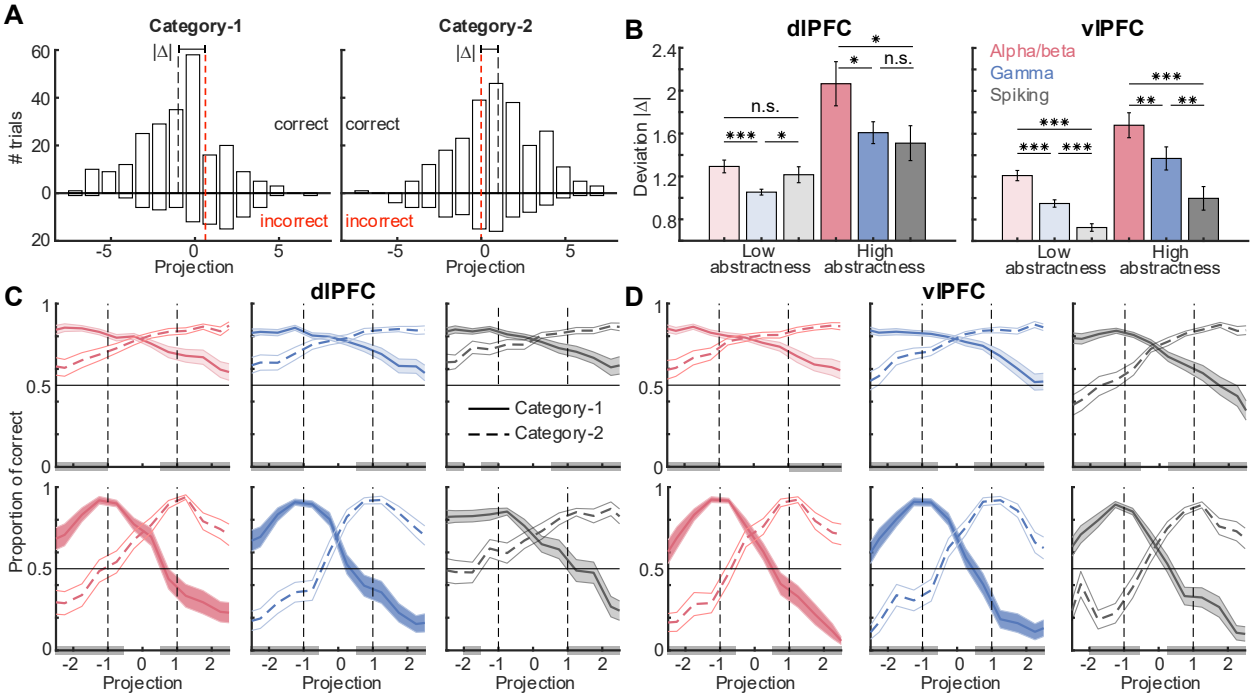

Figure 7. Alpha/beta patterns reflect category errors and predict trial-by-trial decisions.

(A). Histogram of single-trial projections of dlPFC alpha/beta activity onto the category axis for correct (upward bars) and incorrect (downward bars) trials in an example session, pooled across low- and high-abstractness conditions.
*(B). Deviation (|Δ|) between mean projections of correct and incorrect trials under low- and high-abstractness conditions, shown for alpha/beta, gamma, and spiking in dlPFC (left) and vlPFC (right). *** p<0.001, n.s.: not significant, one-sample t-test for each signal type and Wilcoxon rank-sum test for across signals. Error bar: mean ± SE, n = 51.*
(C). Correct choice probability as a function of projection values along the category axis for alpha/beta (left), gamma (middle), and spiking (right) in dlPFC. Solid line: trials where Category-1 is the correct category; dashed line: Category-2 is the correct category. Grey bars on the x-axis represent significant differences between the choice probability of category-1 and category-2 (p<0.001, Wilcoxon rank-sum test). Error bar: mean ± SE, n = 51.
(D). Same as (C) but for vlPFC.

We found the predicted shift in alpha/beta for error trials in which the incorrect category was chosen. **Figure 7A** shows an example session in which alpha/beta projections for incorrect decisions on Category-1 trials shifted toward the Category-2 distribution, and vice versa. We quantified this effect by computing |Δ|, the absolute deviation between the mean projections of correct and incorrect category decisions (see **Methods**). Larger |Δ| values reflect a stronger shift in population activity on error trials, on average. We applied the same analysis to gamma and spiking (**Figure 7B**). We found the average |Δ| deviations across sessions differed significantly from zero (one-sample t-test, p<0.001) for all signals under both low- and high-abstractness conditions. However, alpha/beta projections showed the greatest deviations |Δ| under high-abstractness conditions as compared to gamma and spiking, in both dlPFC and vlPFC. This suggests that while all signals reflect categorization errors, the errors made under higher cognitive demands were most strongly associated with a shift in alpha/beta toward patterns reflecting the incorrect category. This underscores the critical role of alpha/beta in representing context used to guide decisions.

We also found a strong correlation with behavioral performance on a trial-by-trial basis. **Figure 7C** and **7D** shows the proportion of correct category decisions as a function of the single-trial population activity projections. Trials with projections closer to the ideal value for the given category (*x*-axis values of –1 for Category-1, +1 for Category-2) resulted in the highest probabilities of correct decisions. Trials with projections toward the opposite end of the category axis (*x*-axis values of +1 for Category-1, –1 for Category-2) resulted in more incorrect decisions. In between these values, choice probability varied in a graded manner. The slopes of the choice-probability curves were similar across all signal types. This likely reflects the fact that, although alpha/beta showed a larger mean separation between correct and incorrect decisions, it also had greater trial-to-trial variability. In contrast, spiking and gamma showed smaller mean shifts but also lower variance (**Figure S7**), resulting in comparable overall predictive performance across signals.

All together, these results support Prediction 5. Alpha/beta patterns reflect sequence context errors. They reflected category errors. They reflected the NHPs' trial-by-trial category decisions. We also observed that spiking and gamma signals correlated with behavioral performance. This is consistent with the theory: if alpha/beta oscillations

regulate where content-related spiking can emerge, then behaviorally relevant information should also be reflected in spiking and associated gamma activity.

Discussion

Our results provide evidence for the central predictions of Spatial Computing Theory across multiple cognitive tasks. Alpha/beta power conveyed minimal content information (i.e., the items held in WM) but robustly carried information about task context (i.e., task rules, order, and categories). Alpha/beta power carried stronger context information with greater cognitive demands. By contrast, neuronal spiking activity encoded both context and content (**Predictions 1 and 2**). Alpha/beta spatial patterns distinguished between task contexts. This pattern differentiation became greater with greater cognitive demands (**Prediction 3**). Content information carried by spiking was modulated by alpha/beta patterns. Recording sites with lower alpha/beta power had significantly stronger content information. This negative spatial correlation emphasized alpha/beta's role in spatially organizing information processing in the cortex (**Prediction 4)**. Lastly, alpha/beta power closely tracked NHPs' behavioral performance. Alpha/beta power reflected misattributions of task context, whereas content information remained unaffected. Alpha/beta patterns also correlated with NHPs' trial-by-trial decisions about task context, especially under higher cognitive demands (**Prediction 5**). Collectively, these findings validate predictions of Spatial Computing Theory (*1*), demonstrating that alpha/beta oscillations organize cortical activity spatially to support context-dependent cognitive processing.

More broadly, our results are consistent with prior work showing that alpha/beta-band power and coherence in PFC reflects task rules and rule shifts (*20*, *21*). Our observations of increased context information in alpha/beta with increased cognitive demands parallels results from Wutz, et al., 2018 (*13*), who reported increased dlPFC alpha/beta power under similar conditions. Our findings go further by directly quantifying the amount of context information in alpha/beta and showing how it varies with task demands. We also observed greater alpha/beta top-down effects in dlPFC and bottom-up content information in spiking in the vlPFC. This is consistent with prior studies showing that dlPFC is more involved in maintaining abstract rules and goals, whereas vlPFC is more engaged in stimulus-driven encoding and maintenance of object features (*13*, *16–18*). Previous studies have linked alpha/beta oscillations with behavioral performance (*22–24*), as we have in this study. We found that behavioral correlates in the alpha/beta signals were comparably strong to those observed in spiking activity. This suggests that the spatial patterns observed in alpha/beta oscillations are not merely reflective of internal state, but likely play an active role in shaping the neural processes that lead to behavior. They actively structure the neural landscape in ways that influence downstream decision-making.

Our results support the view that alpha/beta oscillations exert inhibitory control over local neural populations (*25*, *26*). Prior studies have shown that alpha/beta power decreases in cortical areas engaged by a task and increases in disengaged areas (*27–29*). In macaque cortex, increased alpha/beta power has been associated with reduced local spiking activity, suggesting a suppressive gating function (*24*, *30–33*). Consistent with

this, we observed a strong inverse relationship between alpha/beta power and content-related spiking.

One possible neural circuit mechanism underlying this spatial regulation involves local inhibitory interneurons. Recent evidence, both experimental (*34–37*) and computational (*38–40*), suggests that alpha/beta oscillations may be generated by somatostatin-expressing (SST) interneurons. These interneurons are well-positioned to mediate top-down modulation (*41*, *42*). Top-down inputs could locally recruit SST cells, leading to an increase in the local alpha/beta power and targeted suppression of nearby excitatory neurons and spiking. In this way, the spatial structure of alpha/beta activity could emerge from selective activation of inhibitory circuits, which in turn sculpt where content-related spiking is permitted or suppressed.

Our findings build on prior work showing that working memory is implemented through interactions between alpha/beta oscillations and brief spiking/gamma bursts (*15*, *33*). More recently, Lundqvist et al. (2023) (*1*) formalized these ideas into Spatial Computing theory, proposing that cortical space is used as a computational dimension for cognitive control. They showed that low-dimensional oscillatory activity carries control-related signals. While that study established key principles of Spatial Computing, several important questions remained open. This includes how spatial patterns of neural activity reorganize with context and scale with cognitive demand, and whether these dynamics relate to behavior. Our results extend this framework by providing direct empirical tests of these predictions. We show that alpha/beta oscillations not only gate information in time but also sculpt the spatial landscape of cortical activity, determining where spiking-based content representations emerge. This offers strong support for the idea that cortical space functions as an active computational resource, with alpha/beta rhythms regulating both the location and timing of information expression.

This spatial regulation offers a mechanistic basis for how the brain flexibly reconfigures its representational geometry with task demands. This is closely aligned with theories of subspace coding (*3*, *4*) and nonlinear mixed selectivity (*5–8*). Subspace coding proposes that different task contexts activate orthogonal patterns of population activity, minimizing interference between representations. Mixed selectivity refers to neurons that respond to specific combinations of stimuli and task context. It is essential for flexible cognition (*6*). While these properties have been observed across the brain, the underlying mechanism has remained unclear. Our results suggest that alpha/beta oscillations provide that mechanism by dynamically regulating where in cortical space certain subspaces and selective combinations can emerge. By spatially suppressing spiking activity, alpha/beta patterns may reduce the overlap between neuronal populations recruited in different contexts, leading to more orthogonal population subspaces. This selective recruitment could also support the emergence of nonlinear mixed selectivity, as neurons participate in different combinations depending on where alpha/beta inhibition is released. In doing so, Spatial Computing explains how a single cortical circuit can express a wide range of task-dependent representations without structural rewiring.

Gamma is typically associated with bottom-up sensory processing and closely tied to local spiking activity (*43–45*). We found gamma carried some content information, though to a much lesser extent than spiking. This likely reflects its lower spatial and temporal resolution relative to spiking activity. Gamma also carried some context information and was correlated with behavioral choices. These findings suggest that while gamma is not traditionally considered a carrier of top-down signals, it may inherit context-related structure through its close coupling with spiking activity that is spatially regulated by alpha/beta oscillations. In this sense, gamma may act as an intermediary, reflecting both the bottom-up content and the spatial constraints imposed by top-down control. Prior work has shown that gamma synchronization is modulated by attention and correlates with behavioral measures, including perceptual accuracy and reaction time (*46*, *47*).

Taken together, our results highlight Spatial Computing as a framework for understanding how the brain flexibly controls cognitive processes. By showing that alpha/beta oscillations dynamically shape when and where content can be expressed, we report a spatial mechanism for cognitive control. This shifts the focus from fixed anatomical circuits to dynamic, spatially structured population activity as the substrate of cognition. Spatial Computing offers a unifying principle that explains how diverse cognitive operations, such as working memory, rule application, categorization, and decision-making, can emerge from a shared circuitry guided by oscillatory control. More broadly, it suggests that cortical space itself is a computational dimension, repurposed in real time to serve changing behavioral goals. This view opens new avenues for understanding the flexible and adaptive nature of neural computation, and for building models of cognition grounded in the analog dynamics of brain activity.

Methods

Non-human primate (NHP) subjects

We analyzed data from three previous studies. The Seq-image Task was from (*9*, *14*), the Single-image Task was from (*1*, *10*, *11*), and the Category Task was from (*12*, *13*, *48*–*50*). For details of experimental design and data collection, please see those publications.

Briefly, each study involved two Rhesus macaques (*Macaca mulatta*) that were trained on a behavioral task using positive rewards (juice), until their performance indicated a solid understanding of all the task requirements. The Seq-image Task included NHPs S (female, 6 years, 6.5 kg) and A (male, 4 years, 10 kg). The Single-image Task included NHPs T (male, 17 years, 13 kg) and D (male, 8 years, 10 kg). The Category Task included NHPs P (female, 8 years, 9 kg) and G (male, 9 years, 13 kg). All NHPs were experimentally naive prior to the experiments analyzed here, except for NHP T, who had previously participated in studies of categorization (*48*) and working memory (*50*); NHP P, who had participated in a previous study of category learning (*51*); and NHP G, who had participated in a study of rule-based categorization (*49*). NHPs were pair-housed, on a 12-hr day/night cycle, and in a temperature controlled environment (80° F). All procedures followed the guidelines of the Massachusetts Institute of Technology Committee on Animal Care and the National Institutes of Health.

Datasets analyzed

For the Seq-image Task, NHPs were trained on two versions of a sequence working memory task. Blocks of each task version—Recognition and Recall—were interleaved, 100–250 trials per block. In both versions, a sequence of two sample images was presented (500 ms each, 1 s blank delay between them). The identity and order of the images then had to be maintained in working memory over a brief delay (1 s). In the Recognition version, a test sequence of two images was then presented, which required a bar-release response during the presentation of the second image if the test images matched the samples in both identity and order. Otherwise, subjects withheld their response until a subsequent matching sequence was shown. In the Recall version, an array of three test images was shown simultaneously, and the subject was required to saccade to the two images that matched the samples in the same order as initially presented. There was no explicit cue to indicate that the task had switched, as it was obvious from context during the first trial after the block switched. NHPs typically performed 2–4 blocks of each task during each recording session. Performance was well above chance on both versions of the task. An infrared-based eye-tracking system (ISCAN) continuously monitored eye position at 60 kHz. In each session, eight tungsten microelectrodes (FHC) were lowered through custom grids using custom-made independently moveable microdrives (Crist). LFPs were referenced to ground, band-pass filtered, and recorded at 1 kHz (Plexon Multichannel Acquisition Processor). Spikes were detected from high-frequency data using an online manually-set amplitude threshold, and manually sorted into isolated single units offline (Offline Sorter, Plexon). For NHP S, a DC-coupled headstage, LFP filtering from 3.3–88 Hz, and spike filtering from 154–8800 Hz was used. For NHP A, an AC-coupled headstage, LFP filtering from 0.7–170 Hz, and spike filtering from 250–8000 Hz was used.

For the Single-image Task, NHPs were trained on a delayed match-to-sample WM task. Each trial began with fixation for 500 ms, followed by a sample object, blank memory delay, and then a free-viewing two-choice test. During the test period, NHPs were required to saccade to the remembered sample and hold fixation there. An infrared-based eye-tracking system (Eyelink 1000 Plus, SR-Research, Ontario, CA) continuously monitored eye position at 1 kHz. Neural signals were recorded with four chronically-implanted 64-electrode (8x8) "Utah" arrays (1.0 mm length, 400 µm spacing, iridium-oxide tips; Blackrock Microsystems, Salt Lake City, UT). Arrays were implanted in both hemispheres of dorsolateral PFC and ventrolateral PFC. LFPs were referenced to ground, AC-coupled, band-pass filtered (0.05–300 Hz), and recorded at 1 kHz (Blackrock Cerebus). Spikes were detected from high-frequency data (250–5000 Hz) using a manually-set amplitude threshold and spike-sorted online using time-amplitude windows.

For the Category Task, NHPs were trained to categorize two categories of dot-pattern category exemplar stimuli, each of which was defined by random distortions around a prototype. Each prototype consisted of 7 pseudo-randomly placed dots. The exemplars were selected by randomly sampling locations from a normal distribution centered at each prototype dot location. The session was split into blocks of trials, with exponentially increasing numbers of category exemplars and increasing levels of distortion around each prototype in successive blocks. The first block featured two exemplars of each category, and each subsequent block doubled the number of exemplars, cumulatively keeping all exemplars from previous blocks and adding an equal number of new exemplars. Each block, the variance of the normal distributions around the prototype dots, which was used to generate each exemplar (i.e. the distortion), progressively increased. NHPs moved from one block to the next when a learning criterion was achieved. Neural signals were recorded with two chronically-implanted 64-electrode (8x8) "Utah" arrays (1.0 mm length, 400 µm spacing, iridium-oxide tips; Blackrock Microsystems, Salt Lake City, UT). Arrays were implanted unilaterally in dorsolateral PFC and ventrolateral PFC. LFPs were referenced to ground, AC-coupled, band-pass filtered (0.75–300 Hz), and recorded at 1 kHz (Blackrock Cerebus). Multi-unit spiking activity was obtained by offline high-pass filtering and thresholding of the 30 kHz-sampled raw broadband signals**.** A channel was defined as recording spikes (multi-unit (MUA) channel) when its spike rate averaged over trials and sample presentation period (0–1 sec) was larger than 1 Hz.

Signal processing

In all datasets, local field potentials (LFPs) were first preprocessed by removing the DC offset, and removing the 60 Hz line noise and its harmonics using an adaptive sinusoid fit method. Subsequently, signals were band-pass filtered between 0.5 and 300 Hz.

For spectral analyses of the LFPs, we applied Morlet wavelet analysis (*52*) after subtracting the evoked potentials phase-locked to trial events, estimated using the grand mean signal across all trials. The Morlet wavelets had a fixed width of 6 cycles, with center frequencies ranging from 2^1 to $2^{7.5}$ Hz in increments of 0.2 (on a logarithmic scale). Resultant spectrograms were then normalized to correct for the ~1/frequency baseline distribution of power by multiplying power at each frequency by the frequency itself. To

obtain band-specific oscillatory power, we averaged the spectral data within the alpha/beta (10–30 Hz) and gamma (30–150 Hz) frequency bands. Oscillatory power was calculated with a temporal resolution of 1 ms.

Neural spiking data across all datasets were smoothed using a Gaussian kernel with a width of 50 ms. The smoothed spiking activity is referred to as the spike rate, with a temporal resolution of 1 ms. Both LFP band power and spike rate data were baseline-corrected by subtracting the mean activity (averaged across trials and time points) in the time interval preceding stimulus onset: [-1, 0] s relative to sample onset for Seq-image Task and Category Task and [-0.5, 0] s for Single-image Task.

In addition to these standardized preprocessing steps, dataset-specific procedures from the original studies were applied as necessary; further details are provided in those publications. Briefly, for the Seq-image Task, sessions containing high-power, broadband frequency artifacts were excluded from analyses. In the Single-image Task, channels identified as having poor electrode contact or significantly elevated noise amplitudes, based on visual inspection, were removed prior to analysis. Unless specified otherwise, all computations below were performed on a session-by-session basis.

Information quantification using PEV

To quantify the amount of information carried by each neural signal type (alpha/beta power, gamma power, and spike rate), we computed the bias-corrected percentage of explained variance (PEV) (*53*) across trials under different task conditions. For different analyses, this quantification was performed at both the population level and the level of individual electrodes or units.

At the population level, we first defined an optimal one-dimensional readout axis based on all correct trials. At each time point, neural population activity (LFP power or spike rate) formed distinct clusters in a high-dimensional space, where each dimension represented one electrode or unit. Each data point corresponded to the neural response from a single trial, and clusters represented two distinct experimental conditions (i.e., a binary classification). For context information, these conditions included Category-1 vs. Category-2 (Category task), first vs. second presentation order (Seq-image task), or Recognition vs. Recall task versions (Seq-image task). For order information in the Seq-image Task, we held the image shown at both sequence orders constant to reduce the effect of stimulus identity on the variance. Each image shown gave us two separate clusters corresponding to first and second sequence order. For content information (e.g., stimulus identity), each stimulus pair formed two separate clusters.

For each binary contrast, we computed the linear discriminant axis, which reflects the optimal linear decoder that maximizes the separation between classes (*54*, *55*) and the Fisher information (*56*). The optimal axis was calculated as the product of the pooled inverse covariance matrix (Σ) of the two clusters and the vector connecting their cluster means: $\Sigma^{-1}(\mu_1 - \mu_2)$. After averaging this optimal axis across time (and across stimulus pairs for content, or across the images shown for order), we projected the original high-dimensional data onto this one-dimensional axis. This step effectively reduced the

dimensionality of the neural data, facilitating the computation of PEV analogous to analyses conducted at the single electrode/unit level.

For the resultant one-dimensional projected population data, as well as for single electrode/unit data, we applied a one-way ANOVA, grouping trials according to different experimental conditions (e.g., stimulus identity, task categories, task rules). Only correct trials were included. A bias correction was applied to PEV estimation (ω^2 formula) to address potential biases associated with small sample sizes (*57*):

$$\omega^2 = \frac{SS_{between\ groups} - df \times MSE}{SS_{total} + MSE}$$

where $SS_{total} = \sum_{i=1}^{N}(x_i - \bar{x})^2$, $SS_{between\ groups} = \sum_{group}^{G} n_{group}(\bar{x}_{group} - \bar{x})^2$, df is the degrees of freedom (i.e., the number of groups G − 1) and MSE is the mean squared error $MSE = \sum_{i=1}^{N}(x_i - \bar{x}_{group})^2$. Consequently, bias-corrected PEV ω^2 provided a robust measure of the amount of information encoded by modulations in neural spiking rates or oscillatory power at both the population and single-electrode/unit levels.

For the Seq-image Task, the content PEV was computed based on the trials of Recog version and trials of Recall version separately for each session. They were aggregated across all the sessions of the two NHPs. The task PEV was computed by all the trials together with the two groups given by “Recog” and “Recall ”. For the order information, we separated each trial into two temporal segments: sample 1 + delay and sample 2 + delay. These two segments from each trial gave the two groups for PEV. The order PEV was computed for Recall trials only so it could be compared with the PEV of order-error trials, which are most easily identifiable in the Recall version.

For the Category Task, the content PEV was computed only over trials from the first block, where the number of exemplar dot patterns was small. This minimized the effects of category learning on the content representation. The category information was computed for the blocks later than block three, by which NHPs typically learned the category (*12*).

To estimate the information during error trials, we projected incorrect single-trial neural responses onto the optimal one-dimensional axis derived from the correct trials. The same one-way ANOVA procedure described above was subsequently applied to these projections.

Low- and high-abstractness based on exemplar distortion

For the Category Task, we classified exemplars into low- and high-abstractness groups using the same criterion as described previously by (*13*). Specifically, the dataset from each recording session was divided into two abstractness levels based on a critical distortion threshold of 1.1 degrees visual angle (DVA), corresponding to the inflection point in behavioral performance. Further methodological details are provided in (*13*). Analysis based on abstractness levels, for example, category PEV, was performed on the trials of low- or high-abstractness separately. For example, for the category information under low- and high-abstractness conditions, we computed the one-dimensional axis from the low- and high-abstractness trials separately and projected the data separately.

Spatial correlation between alpha/beta power and content PEV maps

To quantify the spatial relationship between alpha/beta power and the content-related information (PEV) encoded by gamma power and spiking, we first constructed spatial maps representing both power and content PEV.

For alpha/beta power, we first subtracted the baseline (pre-stimulus) values as described above. Then, for each electrode, power was averaged across all correct trials and across a specific time window during stimulus presentation (approximately 0.1–0.4 s for the Single-image Task and 0.2–0.6 s for the Category Task). This resulted in a vector of 64 alpha/beta power values recorded in each session. These values were then mapped onto the 8-by-8 electrode grid locations to build the power map.

For the content PEV maps, we first computed the PEV of stimulus identity separately for each electrode/unit at each time point. Note that, unlike the population PEV described above, this was computed separately for each individual electrode/unit, without any projection onto a population coding axis. We then averaged the resulting PEV time series within the same time window used for alpha/beta power averaging, producing a content PEV vector with dimension equal to the number of electrodes/units. For gamma activity, the resulting PEV vector was mapped directly to the grid locations of each electrode. For spiking activity, the approach varied slightly: for unsorted multi-unit activity (MUA) recorded in the Category Task, electrode locations were directly mapped. However, for sorted single-unit activity recorded in the Single-image Task, electrodes could have multiple units or none. When multiple units were recorded on the same electrode, their PEV values were averaged to yield a single electrode-level PEV value. Electrodes with no recorded units were excluded from spatial correlation computations.

Spatial correlations were computed using Pearson's correlation coefficient. For each recording session, we correlated the alpha/beta power vector with the corresponding PEV vector, including only electrode locations containing valid data points (i.e., excluding electrodes with missing values because they recorded no units).

Within- and between-category spatial correlation of alpha/beta power maps

For the Category Task, we assessed spatial correlations of alpha/beta power maps both within and between categories. Trials from each recording session were segmented into non-overlapping windows of 50 trials. Within each set, we computed spatial correlations separately and subsequently averaged the results across the trial windows. For each set of 50 trials, we constructed alpha/beta power maps by separately averaging trials from Category-1 and Category-2, both across trials and across the stimulus presentation temporal window. The between-category correlation was calculated directly between these two resulting spatial maps. For within-category correlations, trials from each category were randomly partitioned into two halves, and correlations were computed between spatial maps averaged from these halves. This random partitioning procedure was repeated 20 times, and the average correlation value was taken as the within-category correlation estimate. Finally, spatial correlation values obtained from the trial windows after the third block were averaged to yield the overall within- and between-category correlation estimates for each recording session.

825
826 **Single-trial behavioral performance analysis**
827 Single-trial behavioral performance analyses were conducted following the procedure
828 described by (*19*). For each recording session, neural responses (LFP power or spike
829 rate) were first averaged across the stimulus presentation window for each individual trial,
830 resulting in trial-specific population response vectors. Using only correct trials, we then
831 calculated an optimal one-dimensional discriminant axis by multiplying the inverse
832 covariance matrix of two response clusters (corresponding to Category-1 and Category-
833 2) with the vector connecting their cluster means. Single-trial population responses were
834 projected onto this discriminant axis, and the resultant scalar projections were normalized
835 so that projections of -1 and +1 corresponded to the mean responses of correct Category-
836 1 and Category-2 trials, respectively.
837
838 Incorrect trials exhibited a projection distribution shifted toward the opposite category:
839 incorrect Category-1 trials shifted toward +1, while incorrect Category-2 trials shifted
840 toward -1. For each session, we used the one-sided t-test to see whether this shift in the
841 incorrect distribution was significant ($p<0.05$) or not. For those sessions with significant
842 shifts, we quantified these shifts by computing the mean projection values separately for
843 incorrect Category-1 and Category-2 distribution and then calculated the absolute
844 differences (denoted as $|\Delta|$) between these incorrect-trial means and their respective
845 correct means (-1 for Category-1, +1 for Category-2).
846
847 To estimate the probability of correct choices as a function of scalar projection, we binned
848 the projection values from all trials within each recording session and counted the number
849 of correct and incorrect trials per bin. Given that the projections were normalized for each
850 session, bin alignments were consistent across sessions, with -1 representing the
851 average correct Category-1 response and +1 the average correct Category-2 response.
852 Finally, we aggregated these binned trial counts across all sessions and both NHPs.
853
854 **Statistical Analysis**
855 Statistical tests for significance were performed with a two-sided Wilcoxon rank-sum test
856 (ranksum function in MATLAB) when samples were independent and with a two-sided t-
857 test (ttest function in MATLAB) for paired samples or single samples. Correlations
858 between two variables were computed as the Pearson correlation. Statistical significance
859 was defined by a p value <0.05. The statistical details (correlation coefficient, p value,
860 sample size n) are provided in the figures, figure legends, or the text of the Results
861 section. The specific meaning of the sample size n is clarified when used.

References

1. M. Lundqvist, S. L. Brincat, J. Rose, M. R. Warden, T. J. Buschman, E. K. Miller, P. Herman, Working memory control dynamics follow principles of spatial computing. *Nat Commun* **14**, 1429 (2023).

2. M. Lundqvist, E. K. Miller, J. Nordmark, J. Liljefors, P. Herman, Beta: bursts of cognition. *Trends in Cognitive Sciences* **28**, 662–676 (2024).

3. M. F. Panichello, T. J. Buschman, Shared mechanisms underlie the control of working memory and attention. *Nature* **592**, 601–605 (2021).

4. O. G. Sani, H. Abbaspourazad, Y. T. Wong, B. Pesaran, M. M. Shanechi, Modeling behaviorally relevant neural dynamics enabled by preferential subspace identification. *Nat Neurosci* **24**, 140–149 (2021).

5. J. Duncan, E. K. Miller, "Cognitive focus through adaptive neural coding in the primate prefrontal cortex" in *Principles of Frontal Lobe Function* (Oxford University Press, New York, NY, US, 2002), pp. 278–291.

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

7. K. M. Tye, E. K. Miller, F. H. Taschbach, M. K. Benna, M. Rigotti, S. Fusi, Mixed selectivity: Cellular computations for complexity. *Neuron* **112**, 2289–2303 (2024).

8. S. Fusi, E. K. Miller, M. Rigotti, Why neurons mix: high dimensionality for higher cognition. *Current Opinion in Neurobiology* **37**, 66–74 (2016).

9. M. R. Warden, E. K. Miller, Task-Dependent Changes in Short-Term Memory in the Prefrontal Cortex. *J. Neurosci.* **30**, 15801–15810 (2010).

10. S. Bhattacharya, S. L. Brincat, M. Lundqvist, E. K. Miller, Traveling waves in the prefrontal cortex during working memory. *PLoS Comput Biol* **18**, e1009827 (2022).

11. L. Kozachkov, J. Tauber, M. Lundqvist, S. L. Brincat, J.-J. Slotine, E. K. Miller, Robust and brain-like working memory through short-term synaptic plasticity. *PLOS Computational Biology* **18**, e1010776 (2022).

12. R. F. Loonis, S. L. Brincat, E. G. Antzoulatos, E. K. Miller, A Meta-Analysis Suggests Different Neural Correlates for Implicit and Explicit Learning. *Neuron* **96**, 521-534.e7 (2017).

13. A. Wutz, R. Loonis, J. E. Roy, J. A. Donoghue, E. K. Miller, Different Levels of Category Abstraction by Different Dynamics in Different Prefrontal Areas. *Neuron* **97**, 716-726.e8 (2018).

14. M. R. Warden, E. K. Miller, The Representation of Multiple Objects in Prefrontal Neuronal Delay Activity. *Cerebral Cortex* **17**, i41–i50 (2007).

15. M. Lundqvist, P. Herman, M. R. Warden, S. L. Brincat, E. K. Miller, Gamma and beta bursts during working memory readout suggest roles in its volitional control. *Nat Commun* **9**, 394 (2018).

16. M. Petrides, Lateral prefrontal cortex: architectonic and functional organization. *Philosophical Transactions of the Royal Society B: Biological Sciences* **360**, 781–795 (2005).

17. D. G. V. Mitchell, R. A. Rhodes, D. S. Pine, R. J. R. Blair, The contribution of ventrolateral and dorsolateral prefrontal cortex to response reversal. *Behavioural Brain Research* **187**, 80–87 (2008).

18. T. Yamagata, Y. Nakayama, J. Tanji, E. Hoshi, Distinct Information Representation and Processing for Goal-Directed Behavior in the Dorsolateral and Ventrolateral Prefrontal Cortex and the Dorsal Premotor Cortex. *J. Neurosci.* **32**, 12934–12949 (2012).

19. M. R. Cohen, J. H. R. Maunsell, A Neuronal Population Measure of Attention Predicts Behavioral Performance on Individual Trials. *J. Neurosci.* **30**, 15241–15253 (2010).

20. T. J. Buschman, E. L. Denovellis, C. Diogo, D. Bullock, E. K. Miller, Synchronous Oscillatory Neural Ensembles for Rules in the Prefrontal Cortex. *Neuron* **76**, 838–846 (2012).

21. C. G. Richter, R. Coppola, S. L. Bressler, Top-down beta oscillatory signaling conveys behavioral context in early visual cortex. *Sci Rep* **8**, 6991 (2018).

22. A. Bollimunta, Y. Chen, C. E. Schroeder, M. Ding, Neuronal Mechanisms of Cortical Alpha Oscillations in Awake-Behaving Macaques. *J. Neurosci.* **28**, 9976–9988 (2008).

23. M. Siegel, T. H. Donner, A. K. Engel, Spectral fingerprints of large-scale neuronal interactions. *Nat Rev Neurosci* **13**, 121–134 (2012).

24. S. Haegens, V. Nácher, R. Luna, R. Romo, O. Jensen, α-Oscillations in the monkey sensorimotor network influence discrimination performance by rhythmical inhibition of neuronal spiking. *Proceedings of the National Academy of Sciences* **108**, 19377–19382 (2011).

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

26. E. K. Miller, M. Lundqvist, A. M. Bastos, Working Memory 2.0. *Neuron* **100**, 463–475 (2018).

27. D. Jokisch, O. Jensen, Modulation of Gamma and Alpha Activity during a Working Memory Task Engaging the Dorsal or Ventral Stream. *J. Neurosci.* **27**, 3244–3251 (2007).

28. S. Haegens, V. Nácher, A. Hernández, R. Luna, O. Jensen, R. Romo, Beta oscillations in the monkey sensorimotor network reflect somatosensory decision making. *Proceedings of the National Academy of Sciences* **108**, 10708–10713 (2011).

29. F. van Ede, F. de Lange, O. Jensen, E. Maris, Orienting Attention to an Upcoming Tactile Event Involves a Spatially and Temporally Specific Modulation of Sensorimotor Alpha- and Beta-Band Oscillations. *J. Neurosci.* **31**, 2016–2024 (2011).

30. E. Spaak, M. Bonnefond, A. Maier, D. A. Leopold, O. Jensen, Layer-Specific Entrainment of Gamma-Band Neural Activity by the Alpha Rhythm in Monkey Visual Cortex. *Current Biology* **22**, 2313–2318 (2012).

31. T. van Kerkoerle, M. W. Self, B. Dagnino, M.-A. Gariel-Mathis, J. Poort, C. van der Togt, P. R. Roelfsema, Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences* **111**, 14332–14341 (2014).

32. A. M. Bastos, M. Lundqvist, A. S. Waite, N. Kopell, E. K. Miller, Layer and rhythm specificity for predictive routing. *Proceedings of the National Academy of Sciences* **117**, 31459–31469 (2020).

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

34. T. Womelsdorf, T. A. Valiante, N. T. Sahin, K. J. Miller, P. Tiesinga, Dynamic circuit motifs underlying rhythmic gain control, gating and integration. *Nat Neurosci* **17**, 1031–1039 (2014).

35. J. P. Hamm, R. Yuste, Somatostatin Interneurons Control a Key Component of Mismatch Negativity in Mouse Visual Cortex. *Cell Reports* **16**, 597–604 (2016).

36. G. Chen, Y. Zhang, X. Li, X. Zhao, Q. Ye, Y. Lin, H. W. Tao, M. J. Rasch, X. Zhang, Distinct Inhibitory Circuits Orchestrate Cortical *beta* and *gamma* Band Oscillations. *Neuron* **96**, 1403-1418.e6 (2017).

37. A. B. Van Derveer, G. Bastos, A. D. Ferrell, C. G. Gallimore, M. L. Greene, J. T. Holmes, V. Kubricka, J. M. Ross, J. P. Hamm, A Role for Somatostatin-Positive Interneurons in Neuro-Oscillatory and Information Processing Deficits in Schizophrenia. *Schizophrenia Bulletin* **47**, 1385–1398 (2021).

38. B. Lee, D. Shin, S. P. Gross, K.-H. Cho, Combined Positive and Negative Feedback Allows Modulation of Neuronal Oscillation Frequency during Sensory Processing. *Cell Reports* **25**, 1548-1560.e3 (2018).

39. J. W. M. Domhof, P. H. E. Tiesinga, Flexible Frequency Switching in Adult Mouse Visual Cortex Is Mediated by Competition Between Parvalbumin and Somatostatin Expressing Interneurons. *Neural Computation* **33**, 926–966 (2021).

40. S. Zhao, J. Zhou, Y. Zhang, D.-H. Wang, γ And β Band Oscillation in Working Memory Given Sequential or Concurrent Multiple Items: A Spiking Network Model. *eNeuro* **10** (2023).

41. H. Makino, T. Komiyama, Learning enhances the relative impact of top-down processing in the visual cortex. *Nat Neurosci* **18**, 1116–1122 (2015).

42. F. Studer, T. Zeppillo, T. R. Barkat, Top-down inputs are controlled by somatostatin-expressing interneurons during associative learning. bioRxiv [Preprint] (2024). https://doi.org/10.1101/2024.12.18.629098.

43. P. Fries, T. Womelsdorf, R. Oostenveld, R. Desimone, The Effects of Visual Stimulation and Selective Visual Attention on Rhythmic Neuronal Synchronization in Macaque Area V4. *J. Neurosci.* **28**, 4823–4835 (2008).

44. S. Ray, J. H. R. Maunsell, Different Origins of Gamma Rhythm and High-Gamma Activity in Macaque Visual Cortex. *PLOS Biology* **9**, e1000610 (2011).

45. G. Michalareas, J. Vezoli, S. van Pelt, J.-M. Schoffelen, H. Kennedy, P. Fries, Alpha-Beta and Gamma

Rhythms Subserve Feedback and Feedforward Influences among Human Visual Cortical Areas. *Neuron* **89**, 384–397 (2016).

46. T. Womelsdorf, P. Fries, P. P. Mitra, R. Desimone, Gamma-band synchronization in visual cortex predicts speed of change detection. *Nature* **439**, 733–736 (2006).

47. M. Siegel, T. H. Donner, R. Oostenveld, P. Fries, A. K. Engel, Neuronal Synchronization along the Dorsal Visual Pathway Reflects the Focus of Spatial Attention. *Neuron* **60**, 709–719 (2008).

48. J. A. Cromer, J. E. Roy, E. K. Miller, Representation of Multiple, Independent Categories in the Primate Prefrontal Cortex. *Neuron* **66**, 796–807 (2010).

49. E. G. Antzoulatos, E. K. Miller, Synchronous beta rhythms of frontoparietal networks support only behaviorally relevant representations. *eLife* **5**, e17822 (2016).

50. S. L. Brincat, J. A. Donoghue, M. K. Mahnke, S. Kornblith, M. Lundqvist, E. K. Miller, Interhemispheric transfer of working memories. *Neuron* **109**, 1055-1066.e4 (2021).

51. E. G. Antzoulatos, E. K. Miller, Differences between Neural Activity in Prefrontal Cortex and Striatum during Learning of Novel Abstract Categories. *Neuron* **71**, 243–249 (2011).

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

53. S. Olejnik, J. Algina, Generalized Eta and Omega Squared Statistics: Measures of Effect Size for Some Common Research Designs. *Psychological Methods* **8**, 434–447 (2003).

54. Z. Chen, K. Padmanabhan, Top-down feedback enables flexible coding strategies in the olfactory cortex. *Cell Reports* **38**, 110545 (2022).

55. M. R. Whiteway, B. Averbeck, D. A. Butts, A latent variable approach to decoding neural population activity. bioRxiv [Preprint] (2020). https://doi.org/10.1101/2020.01.06.896423.

56. M. Calderini, J.-P. Thivierge, Estimating Fisher discriminant error in a linear integrator model of neural population activity. *The Journal of Mathematical Neuroscience* **11**, 6 (2021).

57. T. J. Buschman, M. Siegel, J. E. Roy, E. K. Miller, Neural substrates of cognitive capacity limitations. *Proceedings of the National Academy of Sciences* **108**, 11252–11255 (2011).

Acknowledgments
We thank Jordan DeFarias and Meredith Mahnke for technical assistance in data collection. We thank all members of the Miller Lab for helpful comments and discussions.

Funding: This study was funded by Office of Naval Research N00014-22-1-2453 (E.K.M.), Freedom Together Foundation (E.K.M.), The Picower Institute for Learning and Memory (E.K.M.), and ERC starting grant Hot-Coal WM 949131 (M.L.). This material is based upon work supported in part by ARO award W911NF-24-1-022.

Author Contributions: M.R.W, R.F.L., S.L.B., and M.L. collected the data. S.L.B. curated the data. M.L. originated the concept of spatial computing. Z.C. and S.L.B. formalized the central predictions of spatial computing. Z.C. analyzed the data to test the predictions. Z.C., S.L.B., and E.K.M. wrote the manuscript.

Competing interests: The authors declare that they have no competing interests.

Data and materials availability: All data needed to evaluate the conclusions in the paper are available upon reasonable request from the Corresponding Author.

Supplementary Materials

Oscillatory Control of Cortical Space as a Computational Dimension

Zhen Chen, Scott L. Brincat, Mikael Lundqvist, Roman Loonis, Melissa R. Warden, Earl K. Miller

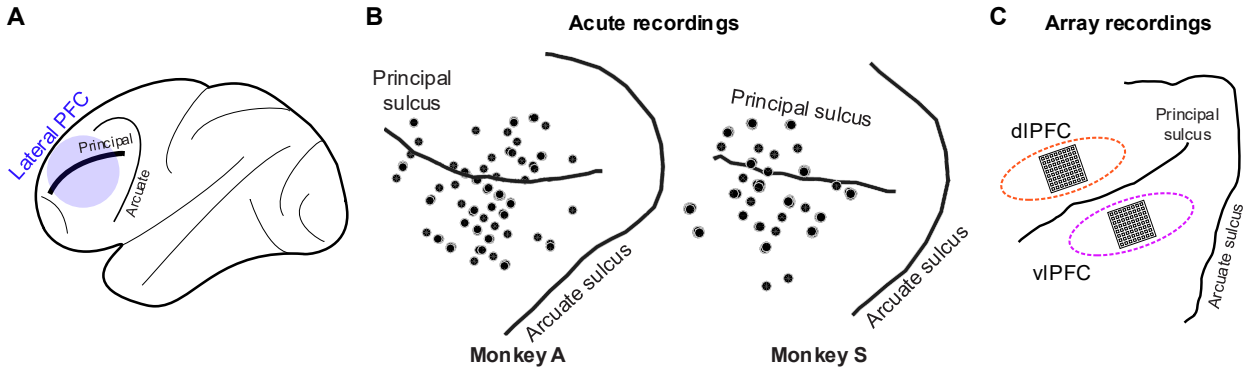

Figure S1. Recording sites in the lateral PFC across tasks

(A). A brain schematic showing the location of lateral prefrontal cortex (PFC) where all the recordings were performed. (B). Anatomical locations of recording sites for the Seq-image Task, for both NHPs. (C). Utah array locations in the dlPFC and vlPFC for the Single-image Task and Category Task.

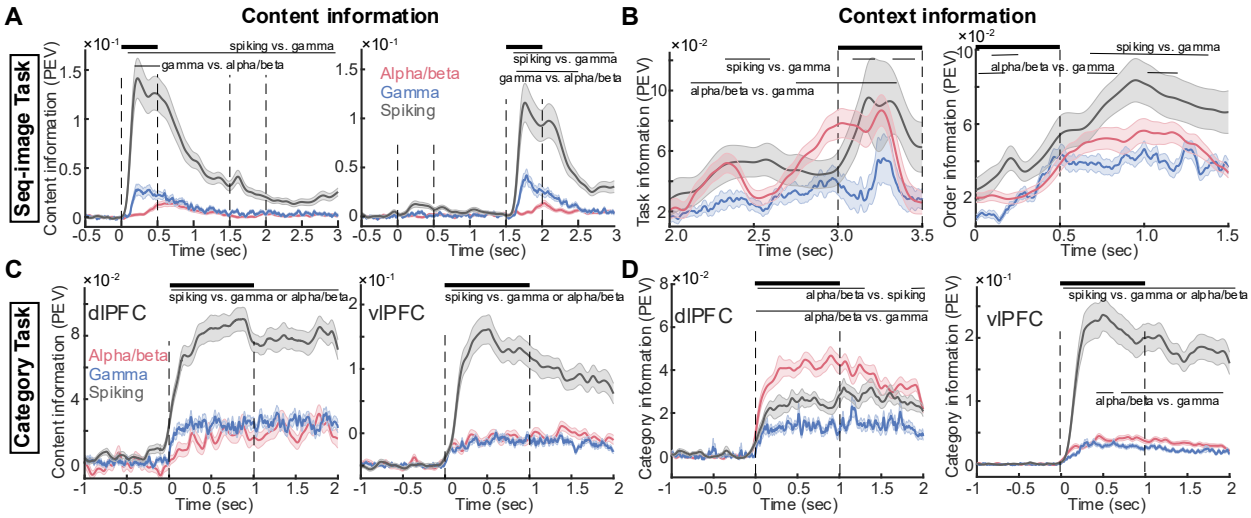

Figure S2. Paired significance test for content and context information

This is the same data (mean ± SE) plotted in main text Fig. 2, but showing markers for significance of pairwise comparisons between signals. (A–B). Content (A) and context (B) information for each signal type in the Seq-image Task. Horizontal black lines on top represent significant PEV values for pairwise comparison between noted signal types ($p<0.001$, Wilcoxon rank-sum). (C–D). Content (C) and context (D) information for each signal type in the Category Task, for dlPFC (left) and vlPFC (right). Horizontal black lines on top represent significant PEV values for pairwise comparison between signal types ($p<0.001$, Wilcoxon rank-sum).

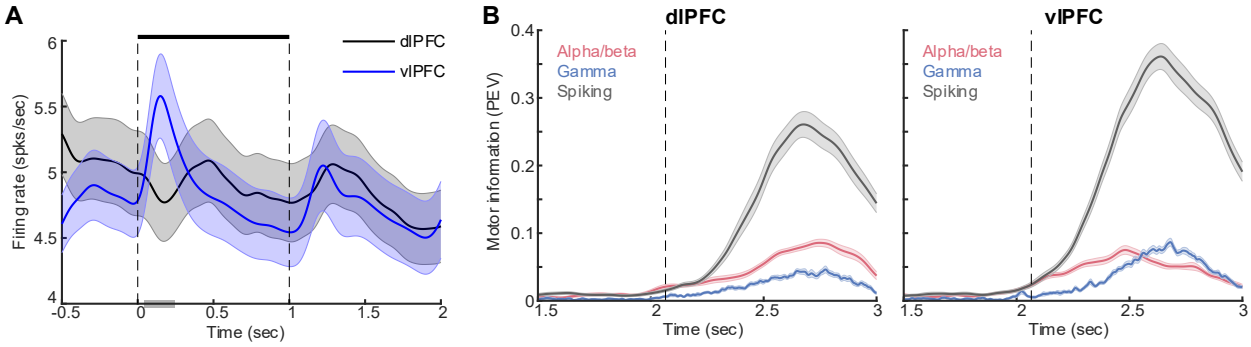

Figure S3. Spike rate and behavioral choice information for Category Task were similar across vlPFC and dlPFC

(A). Spike rate (spikes per second) averaged over all dlPFC and vlPFC MUA channels. (B). Information (PEV) reflecting behavioral choice of the test exemplar on the the left vs. right (indicated by the final gaze location) for dlPFC (left) and vlPFC (right). The vertical dash line indicates the time of test onset. Error bar: mean ± SE, n = 51 sessions.

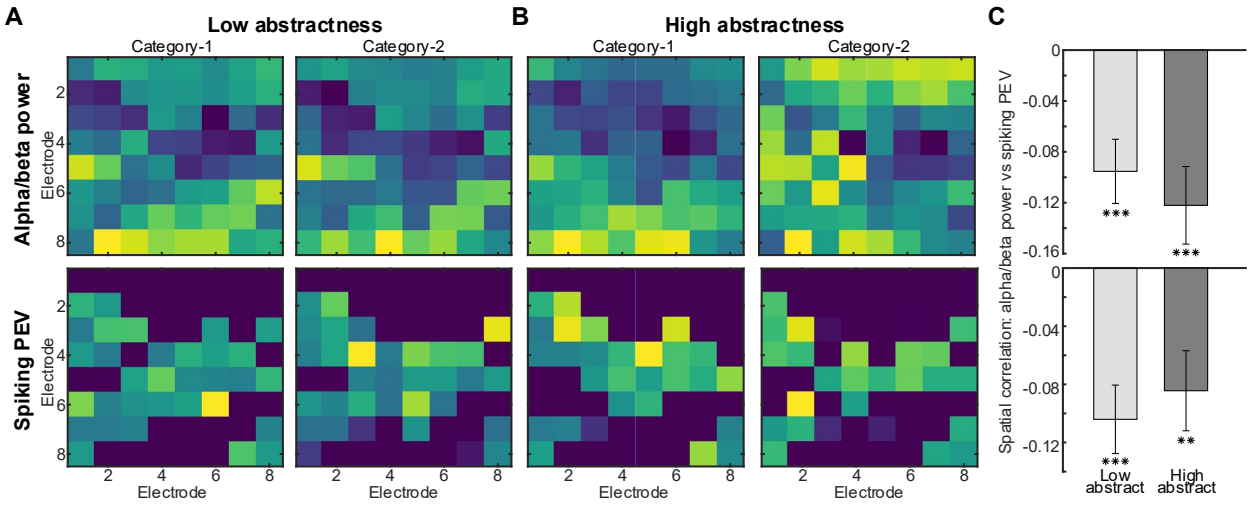

Figure S4. Spatial regulation of content information by alpha/beta patterns for Category Task

(A). Example vlPFC alpha/beta power maps (top), along with the corresponding content PEV carried by spiking (bottom) for both categories under low-abstractness condition. (B). Same as (A) but for the high-abstractness condition. (C). Spatial correlation between alpha/beta power and content PEV in spiking in dlPFC (top) and vlPFC (bottom). *** $p<0.001$, **: $p<0.01$, one-sample t-test. Error bar: mean ± SE. Sample sizes: n = 51.

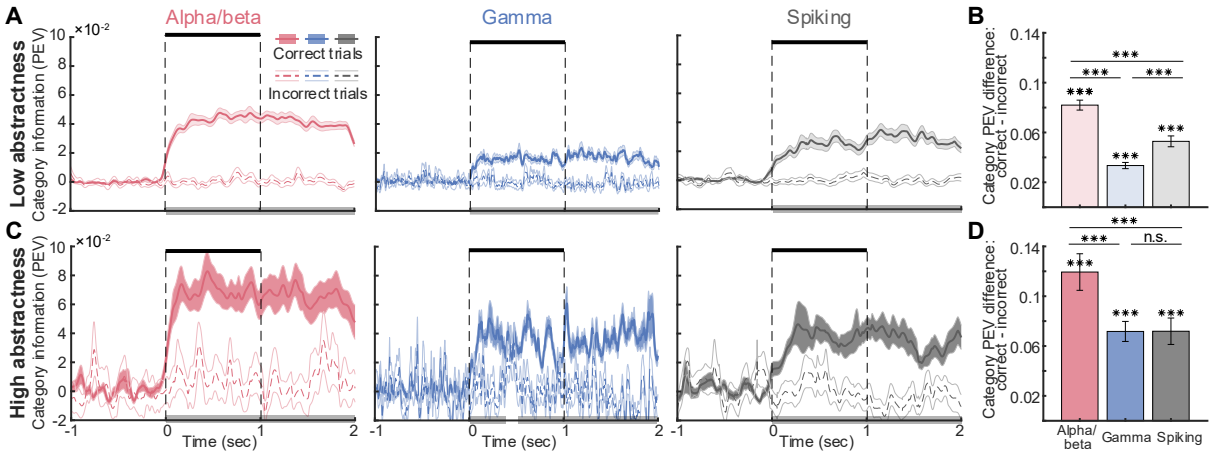

Figure S5. Category information in dlPFC differed between correct and incorrect trials

(A). Category information (PEV) for correct trials and incorrect trials for alpha/beta (left), gamma (middle), and spiking (right) under low-abstractness. Horizontal grey bars on the x-axis represented significant differences between correct and incorrect conditions ($p<0.05$, two-sample t-test). Error bar: mean ± SE, n = 51. (B). Cumulative difference in category information (correct minus incorrect) across time for each signal type. *** $p<0.001$, n.s.: not significant, one-sample t-test for each signal type and two-sample t-test for across signals. Error bar: mean ± SE, n = 51. (C–D). Same as (A–B) but for high-abstractness.

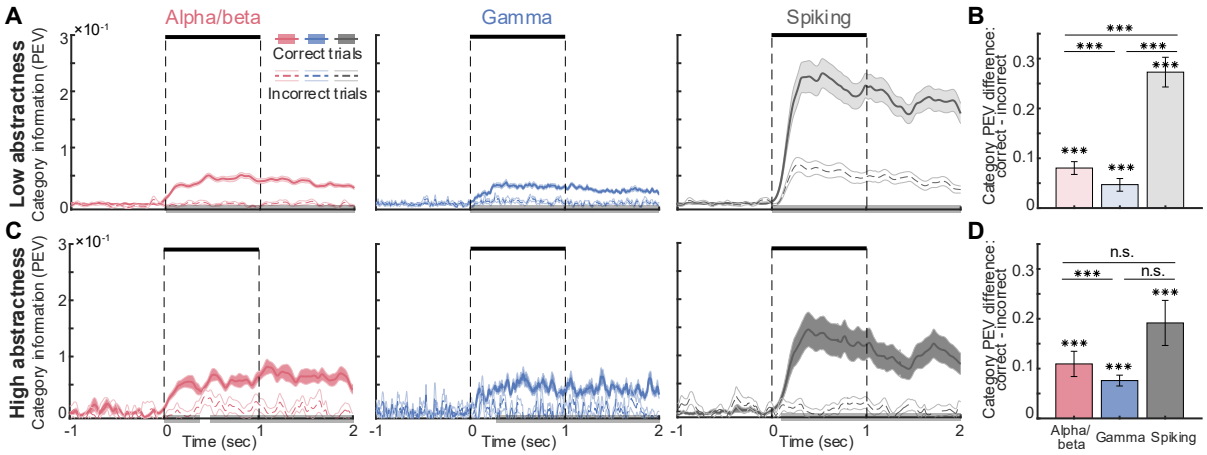

Figure S6. Category information in vlPFC differed between correct and incorrect trials

Same analysis as Fig. S5, but for vlPFC instead of dlPFC. (A). Category information (PEV) for correct trials and incorrect trials for alpha/beta (left), gamma (middle), and spiking (right) under low-abstractness. Horizontal grey bars on the x-axis represented significant PEV differences between correct- and incorrect conditions ($p<0.05$, two-sample t-test). Error bar: mean ± SE, n = 51. (B). Cumulative difference in category information (correct minus incorrect) across time for each signal type. *** $p<0.001$, n.s.: not significant, one-sample t-test for each signal type and two-sample t-test for across signals. Error bar: mean ± SE, n = 51. (C-D). Same as (A-B) but for high-abstractness.

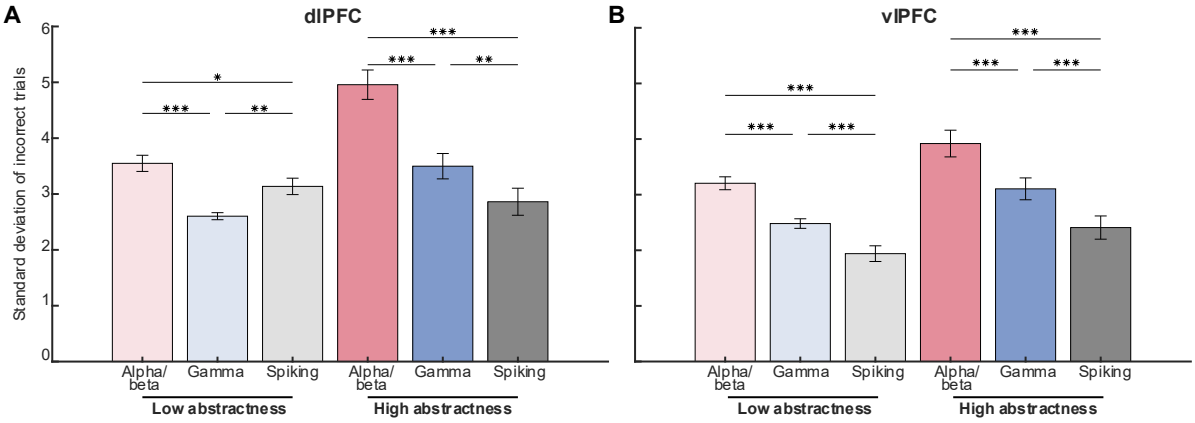

Figure S7. Variability of the projections from incorrect trials for Category Task

(A). Variability (standard deviation) of the incorrect-trial projections for all signal types in dlPFC. *** $p<0.001$, **: $p<0.01$, *: $p<0.05$, one-sample t-test for each signal type and Wilcoxon rank-sum test for across signals. Error bar: mean ± SE, n = 51. (B). Same as (A) but for vlPFC. Alpha/beta showed greater mean differences between correct and error trials than other signals (Fig. 7B), but also greater variability. This resulted in similar overall trial-by-trial correlations with behavior between signal types (Fig. 7C,D).