Prefrontal attentional saccades explore space at an alpha rhythm

Gaillard Corentin[1], Sameh Ben Hadj Hassen[1], Fabio Di Bello[1], Yann Bihan-Poudec[1], Rufin VanRullen[2], Suliann Ben Hamed[1]

1. Institut des Sciences Cognitives Marc Jeannerod, CNRS UMR 5229, Université Claude Bernard Lyon I, 67 Boulevard Pinel, 69675 Bron Cedex, France

2. Centre de Recherche Cerveau et Cognition, Université Paul Sabatier, Université de Toulouse, 31062 Toulouse Cedex 9, France

Corresponding author: Suliann Ben Hamed, benhamed@isc.cnrs.fr

Key words: Monkey, Prefrontal Cortex, FEF, Spatial exploration, Attention, Rhythmic, Alpha Oscillations

Abstract

Recent behavioral studies suggest that attention samples space rhythmically (*1-4*). Oscillations in brain activity have been described as a possible mechanism supporting attentional processes (*5*, *6*). However, the precise mechanism through which this rhythmic exploration of space is subserved by the prefrontal cortical regions at the source of attention control signals remains unknown. Here, we apply machine learning methods to ongoing monkey prefrontal multi-unit population activity, to decode in real-time the (x,y) location of the covert attentional spotlight (*7*), aka the mind's eye. We demonstrate that this prefrontal attentional spotlight continuously explores space at an alpha 7-12 Hz rate. These oscillations determine both neuronal sensory processing, defining how much information is available in the prefrontal cortex about incoming sensory stimuli, and perception, determining whether these incoming sensory stimuli are prone to elicit an overt behavioral response or not. As a result, when sensory events are presented at a specific optimal phase with respect to these oscillations, sensory encoding is reliable and behavior is accurate. When sensory events are presented in anti-phase with respect to this optimal phase, both sensory encoding and behavioral performance are poor. We propose that this rhythmic prefrontal attentional spotlight activity can be viewed as a continuous exploration of space via alpha-clocked attentional saccades. We demonstrate that these attentional saccades are highly flexible, their pattern of space exploration depending both on within-trial and across-task contingencies. These results are discussed in the context of exploration and exploitation strategies and prefrontal top-down attentional control.

Highlights:

- The decoded prefrontal attentional spotlight samples visual space in rhythmic cycles
- This rhythmic attentional exploration predicts neuronal sensory processing accuracy
- This rhythmic attentional exploration predicts overt behavioral accuracy
- These rhythmic cycles define alpha-clocked attentional saccades
- Space exploration by attentional saccades is highly flexible and under top-down control

Introduction

The brain has limited processing capacities and cannot efficiently process the continuous flow of incoming sensory information. Selective attention allows the brain to overcome this limitation by filtering sensory information on the basis of its intrinsic salience (a child crossing the road in front of your car) or its extrinsic value (your old stained coffee mug which you know is somewhere on your crowded desk). Visual selective attention speeds up reaction times (*8*, *9*), enhances perceptual sensitivity and spatial resolution (*10–13*) and distorts spatial representation up to several degrees away from the attended location (*14*). Visual selective attention modulates both neuronal baselines (*15*, *16*) and the strength of visual responses (*17*), decreases neuronal response latencies(*18*), modifies the spatial selectivity profiles of the neurons (*19*, *20*) and decreases shared inter-neuronal noise variability (*21*).

Based on the early work of William James (1890), the spotlight theory of attention assumes that attention is focused at one location of space at a time (*8*, *22*). In this framework, the spotlight is moderately flexible. It is shifted from one location to another, independently from eye position, under the voluntary control of the subject, and its size is adjusted to the region of interest very much like a zoom lens. Converging evidence demonstrate that the prefrontal cortex (PFC) is at the origin of the attentional control signals underlying the behavioral attentional spotlight (*16*, *23–26*). Supporting this idea, we recently demonstrated that this attentional spotlight can be reconstructed and tracked from PFC neuronal population activity with a very high spatial and temporal resolution (*7*, *27*). However, recent experimental work provides a completely different perspective onto selective attention, suggesting that spatial attention samples the visual scene rhythmically[1,2,4,18–23]. These studies report that target detection performance at an attended location fluctuates rhythmically very much like overt sampling processes, such as eye exploration in primates (*28–31*) or whisking in rodents (*32*, *33*). The neural processes at the origin of this rhythmic sampling of space by attention are still poorly understood. Recent works (*5*, *6*) propose that neural oscillations in the fronto-parietal network organize alternating attentional states that in turn modulate perceptual sensitivity.

In the present study, we provide evidence reconciling these two seemingly contradictory views of spatial attention. Specifically, we demonstrate that the 2D (x,y) attentional spotlight decoded PFC activity explores space continuously, through a sequence of attentional saccades that are generated at a specific alpha 7-12 Hz rhythm. Crucially, we show that these oscillations of the attentional locus determine both neuronal sensory processing, defining how much information is available in the prefrontal cortex about incoming sensory stimuli, and perception, determining whether these incoming

sensory stimuli are prone to elicit an overt behavioral response or not. Using Markov chain probabilistic modelling, we further show that space exploration by alpha-clocked attentional saccades depends on both trial and task specific spatial contingencies, implementing an alternation between exploration and exploitation cycles.

Material and methods

Behavioral task and Experimental setup

The task is a 100% validity endogenous cued target detection task (fig 1A). The animals were placed in front of a PC monitor (1920×1200 pixels and a refresh rate of 60 HZ), at a distance of 57 cm, with their heads fixed. The stimuli presentation and behavioral responses were controlled using Presentation (Neurobehavioral systems®, https://www.neurobs.com/). To start a trial, the bar placed in front of the animal's chair had to be held by the monkeys, thus interrupting an infrared beam. The onset of a central blue fixation cross (size 0.7°×0.7°) instructed the monkeys to maintain eye position inside a 2°×2° window, defined around the fixation cross. To avoid the abort of the ongoing trial, fixation had to be maintained throughout trial duration. Eye fixation was controlled thanks to a video eye tracker (Iscan™). Four gray square landmarks (LMs - size 0.5°×0.5°) were displayed, all throughout the trial, at the four corners of a 20°x20° hypothetical square centered onto the fixation cross. Thus, the four LMs (up-right, up-left, down-left, down-right) were placed at the same distance from the center of the screen having an eccentricity of 14° (absolute x- and y-deviation from the center of the screen of 10°). After a variable delay from fixation onset, ranging between 700 – 1200 ms, a small green square (cue - size 0.2°×0.2°) was presented, for 350 ms, close to the fixation cross (at 0.3°) in the direction of one of the LM. Monkeys were rewarded for detecting a subtle change in luminosity of this cued LM. The change in target luminosity occurred unpredictably between 350 – 3300 ms from the cue off time. In order to receive a reward (drop of juice), the monkeys were required to release the bar in a limited time window (150 - 750 ms) after the target onset (Hit trial). In order to make sure that the monkeys did use the cue instruction, on half of the trials, distractors were presented during the cue to target interval. Two types of distractors could be presented: (i) uncued landmark distractor trials (33% of trials with distractor); these corresponded to a change in luminosity, identical to the awaited target luminosity change, and could take place equiprobably at any of the uncued LMs; (ii) workspace distractor trials (67% of trials with distractor); these corresponded to a small square presented randomly in the workspace defined by the

four landmarks. The contrast of the square with respect to the background was the same as the contrast of the target against the LM; when presented at the same radial eccentricity as the LMs, the workspace distractor had the same size as the landmarks; for smaller eccentricities, the size of the workspace distractor was adjusted for cortical magnification such that it activated an equivalent cortical surface at all eccentricities. The monkeys had to ignore all of these distractors. Responding to any of them interrupted the trial. If the response occurred in the same response window as for correct detection trials (150 - 750 ms), the trial was counted as a false alarm (FA) trial. Failing to respond to the target (Miss) similarly aborted the ongoing trial. Overall, data was collected for 19 sessions (M1 10 Sessions, M2 9 Sessions). The behavioral performance of each animal is presented in figure 1B, for hit, miss and false alarm trials. A two-position variant of the above described task was also presented to the monkey. In this task, while the four landmarks were present all throughout the task as previously, only two diagonally opposite positions amongst the four were cued all throughout the session. The pair of cued stimuli changed from one session to the next. 16 such sessions were recorded (8 sessions for M1, 8 sessions for M2). All else was as described for the main four position task.

Electrophysiological recording

Bilateral simultaneous recordings in the two frontal eye fields (FEF) were carried out using two 24 contacts Plexon U-probes (fig. 1B). The contacts had an interspacing distance of 250 µm. Neural data was acquired with the Plexon Omniplex® neuronal data acquisition system. The data was amplified 400 times and digitized at 40,000 Hz. A threshold defining the multi-unit activity (MUA) was applied independently for each recording contact and each session before the actual task-related recordings started.

Neuronal decoding procedure

MUA recorded during the task were aligned on the cue presentation time and sorted according to the monkey's behavioral response (Correct trials, misses trial, false alarm trials). As in Astrand et al. (*7*, *34*), a regularized linear decoder was used to associate, on correct trials, the neuronal activity estimated on a given interval in the cue to target interval and the cued location. The decoder was trained on a random set of 70% of the correct trials at a specific time in the cue to target interval, then tested on the 30% remaining at all time after cue presentation. During training, the input to the classifier was a 48 elements by N-trial matrix corresponding to the average neuronal response on each recording channel for the time interval of interest for each of the N training trials. The imposed output of the classifier was the (x,y) coordinates of the cued landmark for each of these N training trials. During testing, the output of the classifier was estimated for a 48 element vector corresponding to the average neuronal response on

each recording channel for the time interval of interest on a testing trial, new to the classifier. This output can be read as a continuous (x,y) estimate of attention location (*7*) or as a class output, corresponding to one of the four possible visual quadrants (*7*, *34*, *35*). When seeking for a continuous (x,y) readout of attention location, we performed the training using the neuronal activities of Hits averaged over 50 ms immediately before target presentation, then we tested the decoder on neuronal activities averaged over 50 ms all throughout the cue to target interval. When taking a classification perspective, we performed cross-temporal decoding analyses (suppl. figure 1A-B), where successive classifiers were formed based on successive overlapping (every 10ms) time windows during the cue to target interval and tested on independent trials and successive overlapping time windows during the cue to target interval. Mean decoding performance was calculated along the testing axis as the number of correct classifications divided by the total number of classifications. This procedure was repeated 10 times and the grand average over the 10 repeats are used for further analyses. Supplementary figure 1C-H represents this cross-temporal decoding analysis performed onto a training and a testing time interval running from cue presentation to 1200 ms post-cue, when the classifiers are based on neuronal activity sampled over 300, 150, 100, 75, 50 or 25ms. As expected, overall classification performance drops with neuronal sampling window size (*36*). Importantly to the present paper, temporal variations in available content arise at lower sampling window sizes (fig. 2, suppl. Fig. 1F-H). The core analyses of the present paper were performed using a neuronal sampling window size of 50ms.

Oscillations in behavioral performance

Hits and Misses from M1 and M2 were compiled in time (aligned to cue presentation), and merged together across the 19 recording sessions. Behavioral performance, defined as the proportion of (hits/(hits +misses)) was then computed at every millisecond over. The spectral analysis of this time series was performed on detrended data using a Morlet Wavelet transform as in Fiebelkorn et al. (*5*), over the attentional period ranging from 500 ms post cue presentation to 2100ms. Standard error in the power spectrum corresponds to spectral variability during this time interval. Global power spectrum 1/f component was removed from the dataset using a *f normalization (figure 5).

Signal frequency and phase analyses

In the present paper, frequency and phase analyses were performed onto time series (inset in fig. 2A and in fig. 2B) representing attention information classification performance during cue to target interval, for a given training time, along a testing time running from 500 ms to 1200 ms from cue onset. Time series were evaluated at training times ranging from 500 ms to 1200 ms from cue onset, each time series

representing a data sample. Frequency and phase analyses were performed using Wavelet Transform Analyses, based on the Wavelet Coherence Matlab Toolbox (*37*). Specifically, for the time frequency analyses, Morlet wavelet transforms were independently applied to the original data time series (12 Octaves per scale). The significance of peak frequency distributions in the range of interest (7 to 12 Hz) was assessed against the frequency content of time series generated by the random permutation (1000 repetitions, fig. 2B, dashed line) of the MUA time series (prior to decoding). Power to frequency plots are represented with a low frequency cutoff at 4 Hz and normalized by maximal spectrum value. Phase of the signal with respect to cue presentation were obtained from the complex wavelet transform of the signal at the peak frequency of each session.

Characterizing impact of population oscillations onto individual channel spiking activity

For each trial, channel and session, spike trains were smoothed on a 50 ms sliding window over a -700 ms pre-cue to 2000 ms post-cue time series. On the one hand, a Super MUA signal was computed by averaging the spiking activity of the 48 recording channels of each session and each trial. On the other hand, the initial individual channel continuous spiking activity was transformed to identify high-spiking (defined by a spiking rate above 65% of the maximum spiking regime of the individual channel, labelled as 1) and low-spiking (labelled as 0) epochs. The probability of individual channel firing as a function of the oscillatory cycles of the session's Super MUA was then computed as follows. For each channel, for frequencies from 5 Hz to 15 Hz, the spiking probability was computed for the up (+/-π/2 around oscillation peak) and down (+/-π/2 around oscillation trough) oscillatory phases of the frequencies of interest over the entire time window. For each frequency, the analysis time window was adjusted to 1.5 oscillatory cycle length and computations were performed over a minimum of 50 time bins. All further analyses on this metric were performed onto an attentional epoch running from 500 ms post-cue to 2100 ms post-cue.

Peak and trough classification

In order to track whether the frequencies identified on the decoded attentional information causally reflected onto behavior, the following analysis was performed. For each session i, characteristic attention information oscillatory frequency F(i) and Phase P(i) determined using the above described wavelet transform analysis. The decoded classification attention information signal was modeled as a sinusoidal wave determined by the function $MSi(t)=\sin(2.\ \pi.F(i).t-P(i))$. Using this modeled signal (MSi), and based on target time from cue presentation, trials were assigned to one of 10 possible phase intervals ranging from [$-\pi$ $+\pi$] phase offset from the modeled sinusoidal wave For each of these subsets

of trials, decoding accuracy of target location (resp. distractor location) and percentage of hit trials (resp. FA trials) was extracted (fig. 3BC and 4BC). As sensory processing or behavioral outcome could be phase lagged with respect to signal oscillations, target time was progressively shifted using 5 ms steps, so that the phase interval associated with peak sensory processing or behavioral outcome coincided with phase 0. This procedure was applied independently for each of the 18 recording sessions and the outcome of this analysis was then averaged over all sessions, so as to account for variations of F(i) and Phase P(i) from one session to the next. For a precise estimation of phase difference between oscillations in attention information classification decoding and oscillations in sensory processing or behavioral outcome, a circular mean of the corresponding wavelet transform continuous phase difference between the two signals at frequency F(i) was extracted.

Markov chain modeling of spatial attentional exploration strategies

Markov probabilistic chain models were used to characterize the spatial attention exploration strategy of each monkey from cue to target presentation. For each trial, (x,y) time series corresponding to the decoded spatial location of attention during the cue to target interval was collapsed onto the four possible screen quadrants, thus representing how attention moved from one quadrant to the other in time. Based on these discrete time series across all trials of a given session. A Markov chain model was used to estimate the probability that attention stayed in a given quadrant as well as the probabilities that it moved from the current quadrant to one of the three others. This model was performed using the Hmmestimate Matlab function of the Statistics and Machine Learning Toolbox. To compensate for possible idiosyncratic exploration biases of each monkey, the post-cue transition probabilities were normalized with respect to pre-cue spatial attention exploration transition probabilities. Transition probabilities were then normalized for each session and averaged over all sessions and both monkeys. This Markov chain modeling of spatial attentional exploration strategy was independently performed for both tasks: the four cued-location and the two-cued location tasks.

Results

In order to access FEF attentional content in time, we had monkeys perform a cued target-detection task requiring a manual response (fig. 1A) while we recorded the MUA bilaterally from their FEF neuronal ensembles, using two 24-contacts recording probes (fig. 1C). Distractors were presented during the cue-to-target interval and target luminosity was adjusted so as to make the task difficult to perform without orienting attention (fig. 1B). Previous studies demonstrate that PFC based decoding procedures allow to access in which quadrant (*34*, *35*, *38*) or

at which (x,y) location (*7*), attention is placed by the monkeys. In these studies, neuronal signals were averaged over time intervals ranging from 150 ms - 400 ms (*7*, *38*). Larger averaging window sizes produce higher decoding accuracies (suppl. fig. 1), however, larger averaging windows are also expected to result in the smoothing of dynamic changes in the spatial position of attention, artificially reinforcing a static view of the attentional spotlight.

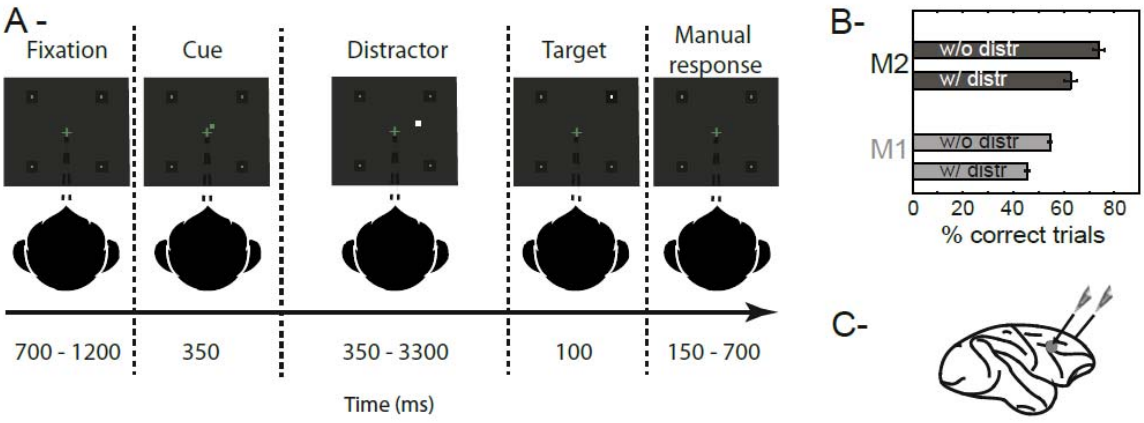

Figure 1: ***Task design and associated behavioral performance.*** (A) 100% validity cued target detection task with temporal distractors. Monkeys needed to hold a bar and fixate a central cross on the screen for a trial to be initiated. The monkeys received a liquid reward for releasing the bar 150 - 750 ms after target presentation. Target location was indicated by a cue (green square, second screen). Monkeys had to ignore any uncued event. (B) Behavioral performance of monkeys M1 and M2 at detecting the target in the presence (w/) or absence (w/o) of a distractor (median % correct +/- median absolute deviations). (C) Recording sites. On each session, 24-contact recording probes were placed in each FEF.

Prefrontal attention-related information oscillates at a 7-to-12 Hz alpha rhythm

In the present study, we seek to characterize spatial attention dynamics in time. As a result, the continuous decoding of attention is performed onto neuronal responses averaged over 50 ms successive (1 ms time steps, suppl. fig. 1) time windows. At this temporal resolution, clear variations in the prefrontal attention-related information are observed. Indeed, when a classifier is trained to decode attention at a given time from cue onset, and tested onto novel activities recorded during the cue to target interval (cross-temporal decoding analysis, fig. 2A). Fluctuations in instantaneous classification accuracies can be noted, at a distance from cue processing. These fluctuations are reliably associated with a distinct peak in the power spectrum relative to chance, in the 7 - 12 Hz range. This is illustrated for an exemplar session in fig. 2B. The power spectrum was quantified using a Morlet wavelet transform analysis (WTA) performed onto independent session time series ranging from 700 – 1200 ms following cue onset (fig. 2B, inset), and assessed against the 95% confidence interval as defined by a random

permutation procedure (see methods, fig. 2B, red line). Peak frequency in the 7 - 12 Hz range is assessed using this method for each monkey and each session individually. For the exemplar session in fig. 2B, peak frequency is identified at 9.2 Hz. Overall, inter-individual and inter-session variability was low and prefrontal attention-related information oscillated, in monkey M1 (resp. M2), at an average frequency of 9 Hz (fig. 2C, resp. 8.6 Hz). A clear phase-locking between these attention-related oscillations and cue onset can be seen across both monkeys (fig. 2D, M1: -75°; M2 -65°). This rhythmic oscillation of the prefrontal attentional spotlight is phase reset by cue presentation and actually pre-exists to cue presentation (see below).

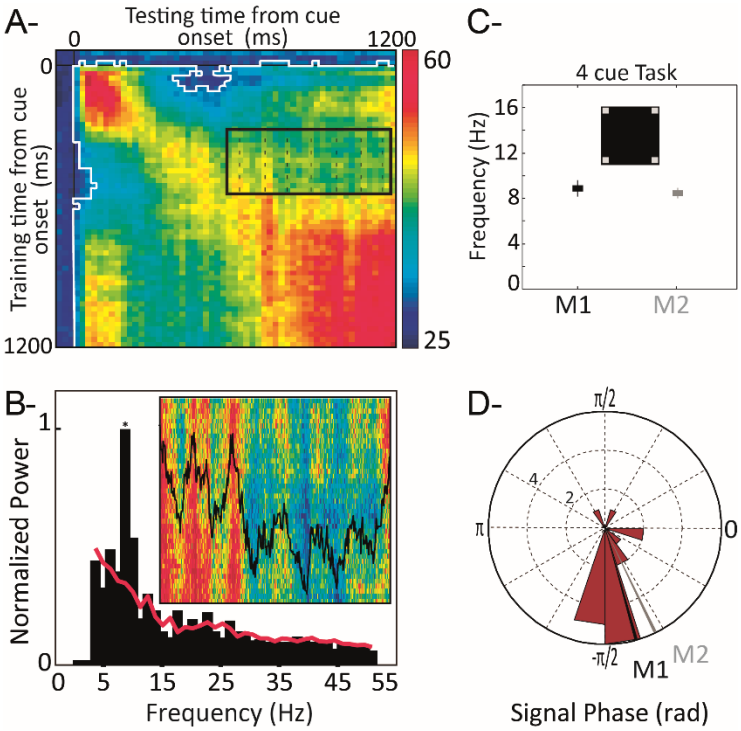

Figure 2: ***Oscillation of prefrontal attention-related information.*** (a) Cross-temporal classification from 100 ms before to 1200 ms after cue presentation (step of 10ms, averaging window of 50ms) for an exemplar session. Each pixel represents mean decoding performance obtained over 1000 repeats of 70% training trials and 30% testing trials). White contour represents the 95% confidence interval as assessed from trial identity random permutation. Black contour represents inset in fig. 2B. (b) Inset: close-up of the cross-temporal classification (500 ms post-cue to 1150 ms post cue along testing time 400 ms post-cue to 575 ms post cue along training time) and corresponding mean classification along testing time (black). Normalized power in this cross-temporal classification interval, as assessed with a wavelet transform analysis (red line: 95% confidence interval) for the exemplar session presented in (a). (c) Average +/- s.d. of peak power in a 7 - 12 Hz interval, over all sessions, for each monkey (M1:

black, M2: gray), in the 4-cued locations version of the task. (d) Circular distribution of signal phase with respect to cue onset, at identified peak frequency (mean phase: M1: black, -75°, M2: gray, -65°).

Alpha rhythm paces FEF population code

Oscillations in the attention-related population activity can either reflect a global rhythmic entrainment of the entire FEF population or changes in the FEF population code at a specific frequency. Fig. 3A represents, for one exemplar recording probe, on an exemplar trial, and for each recording channel, the time epochs at which spiking-rate exceeds the 65% of the maximum spiking regime of the individual channel. The peaks of the alpha oscillations are identified on the super MUA of the same individual trial (*39*) (fig. 3B,) and plotted against the spiking probability changes represented in fig. 3A. The high spiking probability epochs of individual channels coincide with peak alpha oscillatory phases in the super MUA. This is captured by a spectral analysis of changes in spiking probability in a frequency range running from 5 to 15 Hz (see material & methods). Most channels of fig. 3A display a modulation of spiking probability in an 8 - 12 Hz frequency range (fig. 2C, color code matching fig. 2A). This holds true for all sessions (fig. 2D, mean+/-s.e.). However, this alpha rhythmic modulation of spiking probability does not reflect a global entrainment of the entire population. Rather, the channels with highest change in normalized spiking activity change from one super MUA alpha peak to the next, thus reflecting a change in the FEF population code. These variations correspond to changes in the spatial allocation of the attentional spotlight that will be described hereafter.

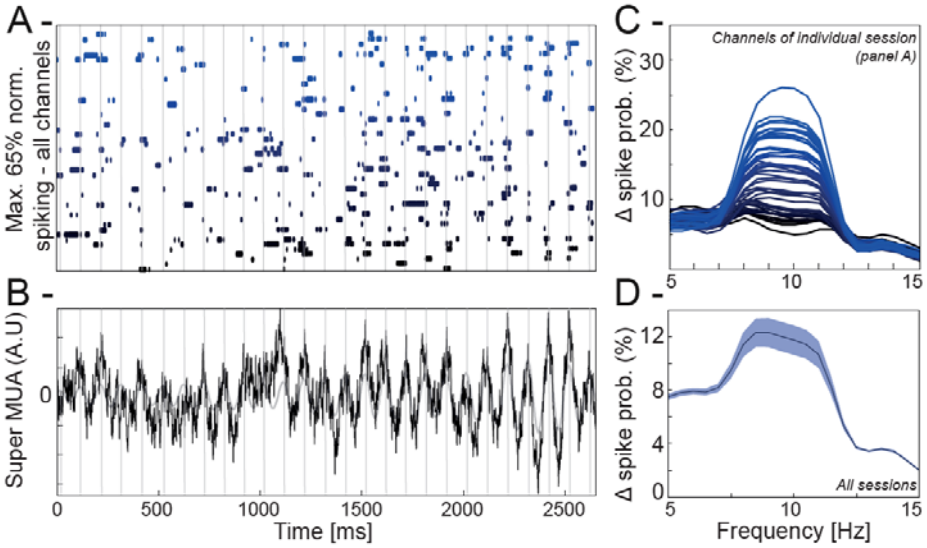

Figure 3: **Alpha rhythm paces FEF population code.** (A) Individual channel spiking probability at a threshold of 65% (1 trial, 48 channels) as a function of time. (Cue is presented at 700ms. Grey vertical lines: peak of alpha cycles of the super MUA in (B). Individual channels are ordered and color coded in a gradient of blue, as a function of the power of their alpha locking presented in (C). (B) Raw and alpha

253 filtered single trial population super MUA calculated over the 48 MUA channels. Grey vertical lines:
254 peak of alpha cycles of the super MUA. (C) Changes in individual channel spiking probability, across all
255 trials, as a function of putative locking to frequencies from 5 to 15 Hz. Spiking probability is specifically
256 affected in the alpha frequency. Channels are color coded in a gradient of blue, as a function of the
257 power of their alpha locking. (D) Mean+/-s.e. phase frequency modulation of spiking activity across all
258 sessions and all channels.

259

260 **Oscillations in prefrontal attention-related information predict prefrontal target encoding and target detection**

261 In order to quantify the relationship between prefrontal attention-related oscillations and both target
262 processing and target detection, trials were classified, for each session, as a function of when the target or the
263 behavioral response were presented relative to the prefrontal attention information oscillation peak (fig. 4A). To
264 this aim, the oscillations in each session were modeled by a sinusoidal wave with the session's specific oscillatory
265 frequency and cue phase-shift. Targets were assigned to phase bins of width of $2\pi/10$, thus covering an entire
266 oscillation cycle.

267 In a first step (fig. 4BC), we focused onto prefrontal target processing. For hit trials, we quantified how
268 much information was available about the target in the prefrontal neuronal population as follows. Neuronal
269 activities were averaged between 50 – 100 ms post-target and used to quantify the accuracy of a four-class
270 classifier at assigning target location to the actual quadrant it was presented in, as compared to the other uncued
271 quadrants. Classifier training and testing were performed onto independent trial subsets from the same category.
272 For each session, target-related prefrontal decoding accuracy was then computed for each independent bin of
273 target-to-attentional oscillation phase-relationship. To increase the resolution of this analysis, this operation was
274 repeated with successive phase bins shifted by 5% of their width. The lag that generated the highest discrimination
275 between maximum and minimum decoding accuracy in the cycle was used to define optimal phase-shift between
276 sensory processing and attention signal oscillations (*18)* (fig. 4B). An average difference in peak and trough
277 decoding accuracies of 10% can be noted when decoding accuracies are cumulated, across all sessions, at optimal
278 phase-shift between sensory processing and signal oscillations (fig. 4B). This difference is highly systematic as
279 illustrated is figure 4C for each session and each monkey independently. The average target decoding accuracy at
280 peak for monkey M1 (resp. monkey M2) was of 54% +/- 4 (resp. 58%+/- 2). At trough, these values dropped to 44%
281 +/- 3 (resp. 47%+/- 1.5). In contrast with the low degree of inter-session variability that we report for prefrontal
282 attention information locking to cue onset (fig. 2D), phase lag between signal and optimal target processing was
283 quite variable (fig. 4B, inset). This variability correlated with intersession behavioral variability in reaction times (fig.
284 5, discussed below). Overall, these results demonstrate a direct modulation of FEF target encoding by the ongoing
285 alpha oscillations that we characterize on the prefrontal attention information.

286

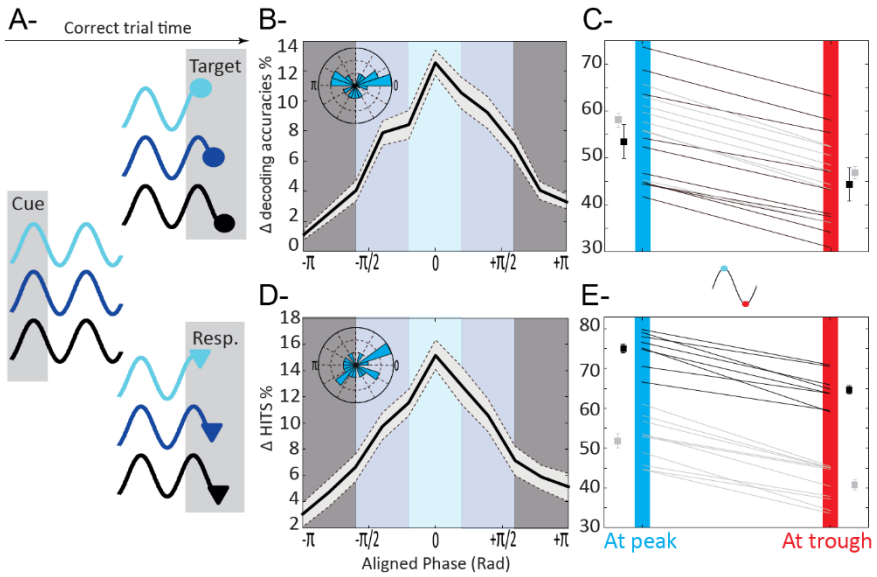

Figure 4: ***Prefrontal target related information and hit rates depend on when the TARGET is presented relative to attentional oscillation cycles.*** (A) Categorization schema of trials as a function of when the target (top) and the behavioral response (bottom) were presented relative to attentional oscillation cycles. (B) Target related information is significantly higher in trials presented at optimal phase with respect to the attentional oscillation cycles. Zero phase corresponds to optimal phase and not to zero phase-locking relative to the attentional signal, hence the observed phase shift distribution (inset: radial distribution of phase shifts relative to cue presentation, 18 sessions). (C) Peak (upper third of the distribution in B-) to trough (lower third of the distribution in B-) variations in target related information, for each monkey (M1: black, M2: gray), for each session. (D) Percentage of hits is significantly higher in trials presented at optimal phase with respect to the attentional oscillation cycles (inset: radial distribution of phase shifts relative to cue presentation, 18 sessions). (E) Peak to trough variations in percentage of hits, for each monkey (M1: black, M2: gray), for each session.

In a second step (fig. 4DE), we used the same procedure as described above, in order to quantify, how target detection (hit rates) depended on target presentation time relative to the prefrontal attention information oscillation cycle. Again, as for target processing, the lag that generated the highest discrimination between maximum and minimum hit rates in the oscillatory cycle was used to define optimal phase-shifts between target detection and signal oscillations (fig. 4D, inset). An average difference in peak and trough decoding accuracies of 10% can be noted when decoding accuracies are cumulated at optimal phase-shift between target detection and signal oscillations (fig. 4D). Again, this difference is highly systematic as illustrated in fig. 4E for each session and each monkey independently. The average target detection at peak for monkey M1 (resp. monkey M2) was of 75 +/- 1.5 (resp. 52%+/- 2). At trough, these values dropped to 64.5% +/- 1.5 (resp. 41.5%+/- 2). As a result, when hit rates

are calculated, across all session and both monkeys, as a function of target presentation time with respect to cue onset (fig. 5A), two significant oscillatory peaks are observed onto behavior, one in the theta (3 to 5 Hz) frequency band, and one in the alpha (9 to 14 Hz) frequency band (fig. 5B), thus reproducing previous behavioral observations (*1*, *2*, *4*, *5*, *40*, *41*). Overall, we show a direct modulation of behavioral target detection by the ongoing alpha oscillations in prefrontal attention information, mirroring our observations onto target processing by prefrontal cortex.

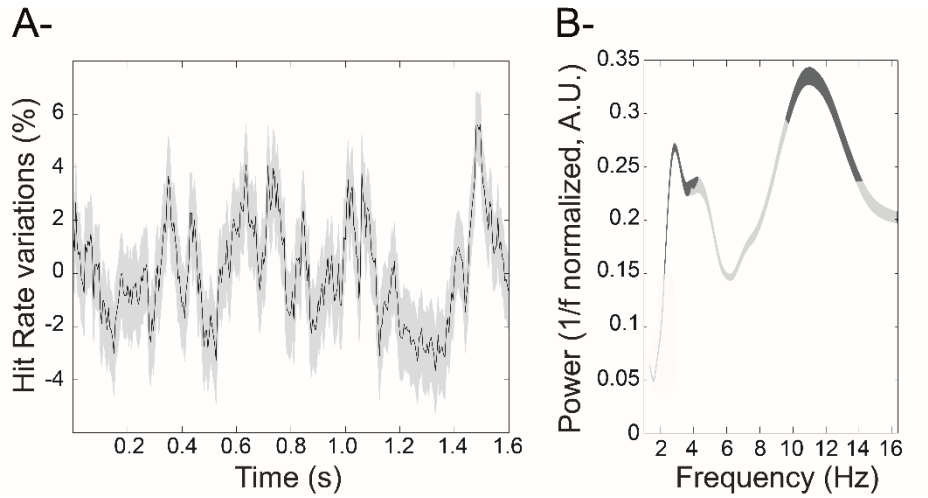

Figure 5: ***Oscillations in cumulated behavioral performance.*** (A) Changes in hit rates (detrended mean +/- s.e.) as a function of time of target presentation relative to cue presentation. Behavioral data compiled across all recording sessions and both monkeys. (B) Complex Morlet Wavelet analysis of behavioral time series to extract power spectrum (mean +/- s.e.). Frequencies significantly modulating overall behavioral performance in dark gray (95%CI on random permutation of target timings).

However, at a closer look, and as reported above for target processing, phase lag between signal and optimal target detection was quite variable from one session to the next (fig. 4D, inset). Phase-lag between optimal target processing and optimal target detection (fig. 6A) was also quite variable (fig. 6B). Importantly, these phase shifts positively correlated with median session reaction times (r^2=0.252, p=0.033). In other words, this phase relationship was predictive of the monkey's response speed in the sessions, and possibly reflects differences in global motor preparation or task engagement states from one session to the next, independently from the ongoing attentional oscillations.

The above reported effects of prefrontal attention information oscillations onto target processing and behavioral outcome can either be interpreted in terms of modulations in attentional focus (i.e. the degree to which

attention is dedicated to sensory processing) or in terms of displacement of the attentional spotlight. In the following, we provide robust evidence in favor of attentional displacement.

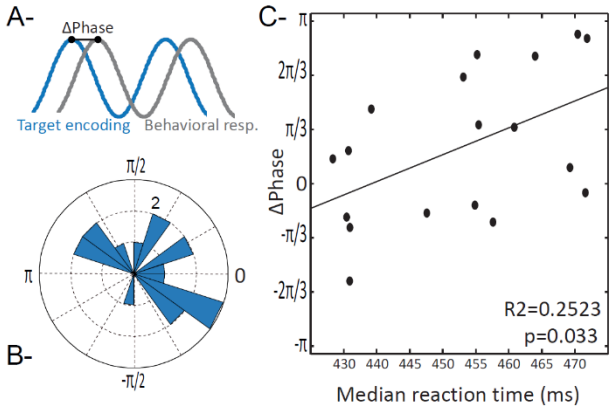

Figure 6: ***Phase lag between optimal target encoding and optimal target detection behavioral response*** (A) vary from one session to the next (B) and are positively correlated with median reaction times on each session (C).

Oscillations in prefrontal attention-related information predict prefrontal distractor encoding and false alarm production

Here, we explore the incidence of the oscillations in prefrontal attention-related information onto the processing of uncued distractors and the production of false alarms (fig. 7). Along the same experimental procedure used in the previous section to explore the incidence of the oscillations in prefrontal attention-related information onto the processing of cued targets and the production of hits. We first focused onto prefrontal distractor representation (fig. 7B). An average difference in peak and trough distractor decoding accuracies from prefrontal neuronal responses of over 30% can be noted when decoding accuracies are cumulated at optimal phase-shift between distractor sensory processing and signal oscillations (fig. 7B). This difference is highly systematic across sessions and monkeys (fig. 7C). The average distractor decoding accuracy at peak for monkey M1 (resp. monkey M2) was of 45% +/- 2 (resp. 43%+/- 1.7). At trough, these values massively dropped to 14% +/- 3.5 (resp. 7%+/- 3). As observed for target processing, phase lag between signal and optimal distractor processing was quite variable (fig. 7B, inset).

In a second step, we quantified how responses to distractors (false alarm rate) depended on distractor presentation time relative to the prefrontal attention information oscillation cycle. An average difference in peak and trough false alarm rate of more than 10% can be noted when false alarms are computed at optimal phase-shift between distractor detection and signal oscillations (fig. 7D). This difference is highly systematic across sessions and monkeys (fig. 7E). The average distractor detection at peak for monkey M1 (resp. monkey M2) was of 45% +/-

1.5 (resp. 42%+/- 2). At trough, these values dropped to 36% +/- 2 (resp. 29%+/- 2.5). As seen for hit rates, phase lag between signal and optimal distractor detection was quite variable (fig. 7D, inset).

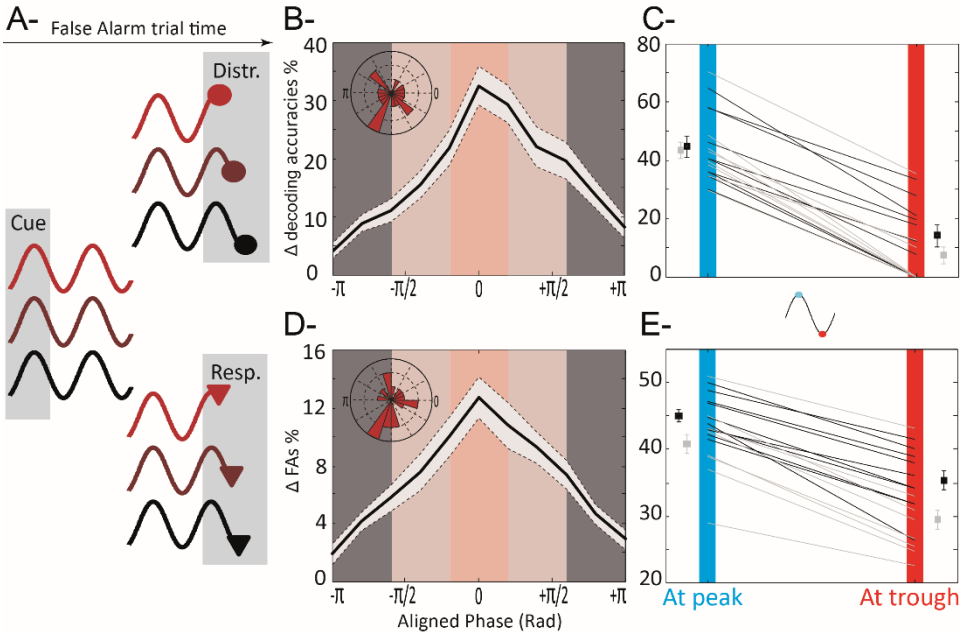

Figure 7: ***Prefrontal distractor related information and false alarm rates depend on when the DISTRACTOR is presented relative to attentional oscillation cycles.*** (A) Categorization schema of trials as a function of when the distractor (top) and behavioral response (bottom) was presented relative to attentional oscillation cycles. (B) Distractor related information is significantly higher in trials presented at optimal phase with respect to the attentional oscillation cycles. (C) Peak to trough variations in distractor related information, for each monkey (M1: black, M2: gray), for each session. (D) Percentage of false alarms is significantly higher in trials presented at optimal phase with respect to the attentional oscillation cycles (inset: radial distribution of phase shifts relative to cue presentation, 18 sessions). (E) Peak to trough variations in percentage of false alarms, for each monkey (M1: black, M2: gray), for each session. All else as in fig. 3.

Overall, we thus show a direct modulation of how the PFC represents distractors as well as the overt behavioral responses to distractors by the ongoing prefrontal attention information alpha oscillations, mirroring our observations onto target processing and target detection. These observations support the hypothesis of a displacement of attention in space. In the following, we provide evidence for an explicit link between the above described oscillations in prefrontal attention information and exploration of space by a highly dynamic and rhythmic attentional spotlight operating in the alpha frequency range.

Attentional "saccades"

The above described analyses were performed on a quantification of the accuracy with which attention could be localized in one of the four visual quadrants, based on the observed prefrontal population neuronal response. In a previous study (*7*), we demonstrated that the continuous (x,y) readout of a linear classifier assigning neuronal activities to a spatial location of attention is a relevant proxy for a real-time access to the attentional spotlight represented in the PFC. Importantly, this continuous (x,y) readout of the prefrontal attention spotlight is predictive of behavior, both in terms of hit and false alarm rates. In the following, we apply the same approach to extract (x,y) attention spotlight trajectories in time before and after cue presentation, to the major difference that the readout is obtained at higher temporal resolution, from neuronal responses averaged over 50 ms rather than on 150 ms as presented in the Astrand et al. (*7*) Movie 1 presents such prefrontal attention spotlight trajectories for an exemplar trials, during all the time of the trial. The attentional spotlight is not stable, nor is it hopping between the four most salient locations. Rather, it is exploring space through a succession of attentional "saccades" that bring the spotlight from one location in space to another, both around and away from the cue.

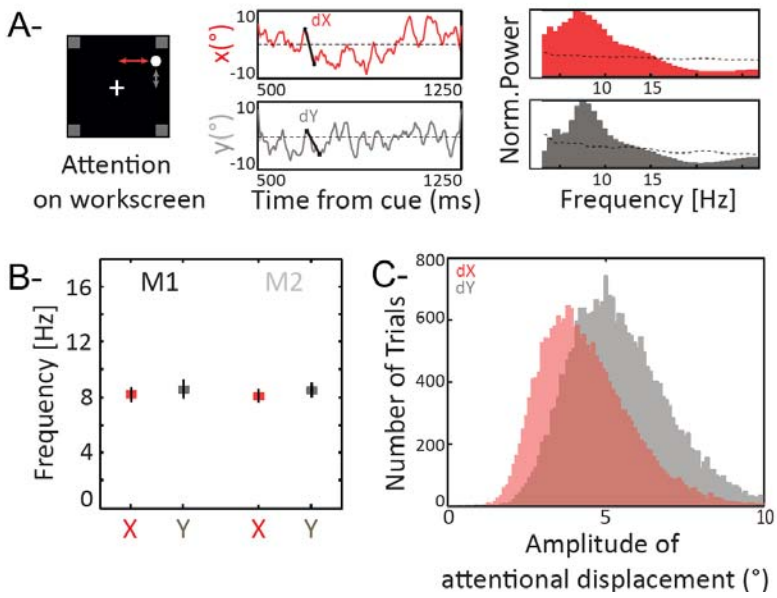

Figure 8: ***The spatial oscillations in prefrontal attentional information map onto changes in attentional spotlight position in space.*** The (x,y) position of the attentional spotlight (A, left panel: x, red; y, gray) varies in time (A, middle panel) at a frequency (A, right panel, dotted black line, 95% confidence interval) that matches the frequency characterized in overall prefrontal attentional information (cf. figure 2). (B) Mean maximal power peak frequency of attentional spotlight trajectory

for M1 and M2, in x (red) and y (gray). (C) Distribution of amplitude of attentional displacement during the cue to target interval (500 to 1250 post cue), along the x (gray) and in y (red) dimensions.

The projections of an exemplar prefrontal attention spotlight trajectory onto the x- and y-dimensions are presented in fig. 8A (middle panel), as well as their spectral power distribution (right panel). A systematic rhythm in attentional displacement can be identified on both x- and y-traces, on all trials and all sessions, for each monkey (fig. 8B, monkey M1: X=8.1 Hz+/ -1.6, Y=8.3 Hz+/-2; monkey M2: X=8 Hz+/-1, Y=8.4 Hz+/-2), in the same range as identified for the global attention population information. No statistical difference is observed between either the peak frequency in the 7-12 Hz range identified in either the x- and y- attentional traces and the peak frequency identified in the global attention information content ($p= 0.49$ and $p= 0.87$ respectively, data not shown), confirming a strong link between these measures. These prefrontal attention spotlight trajectories are exploring space homogenously. Interestingly, a significant difference was observed between the distributions of attentional displacement along the x- and y-axis ($p<0.0001$), indicating a larger exploration of space along the vertical dimension. In this last section, we will demonstrate that while the rhythmic exploration of space by the prefrontal attentional spotlight is both *rhythmic* and *continuous*, how space is being explored is determined by task-related top-down contingencies.

Task-contingencies define where rhythmic attention is deployed in space

During cued target detection tasks, the cue serves to orient attention towards the spatial location where the target is expected to be presented. The absolute distance between two successive attentional saccades does not vary between before (fig. 9A, black) and after cue presentation (fig. 9A, grey, Kolmogorov-Smirnov test $p>0.99$). However, the spatial distribution of these attentional saccades vary significantly between pre-cue and post-cue epochs. Specifically, fig. 9B represents the heat maps of the spatial distribution of the decoded attentional spotlight during the pre-cue interval (-500 to –200 ms, contour 1) and the post-cue interval (500 - 1200ms, contour 2), for each category of cued trials (T1, T2, T3 and T4). During the pre-cue epoch, the heat maps are centered onto the fixation point (median 0.9°+/-0.07°), exploration being confined within the 10.7° central degrees. During the post-cue epoch, the heat maps shift towards the cued landmark by, on average, 3.6° (+/-0.2). For all cued conditions, attentional exploration, extends up to 14.5° towards the cued location (exploration probability threshold of 60%). As a result, rhythmic attention is deployed onto the spatial map changes as a function of within trial task-contingencies.

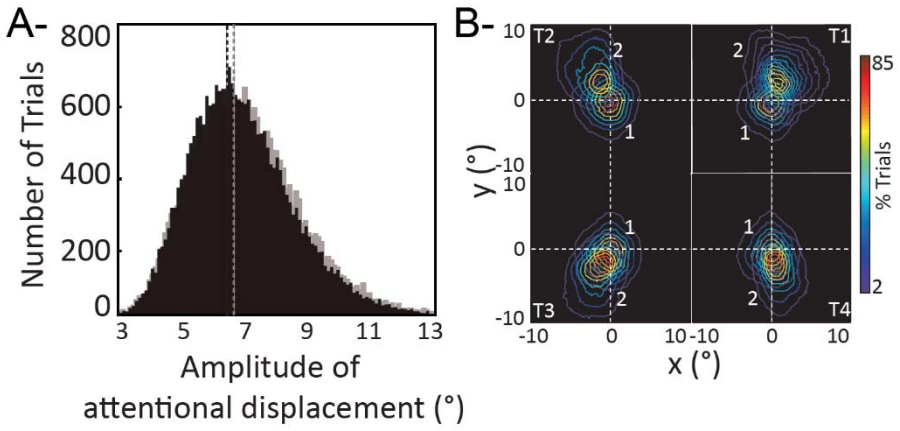

Figure 9: ***Attentional spotlight exploration is an endogenous process affected by task event (CUE).*** (A) Distribution of amplitude of attentional displacement between one prefrontal attentional position and the next, in the pre-cue period (black) and in the cue-to-target interval (gray). (B) Heat maps of the spatial distribution of the decoded attentional spotlight between preCue (1) and postCue (2) epochs.

To further explore the influence of task contingencies onto the spatial deployment of rhythmic attention, we used Markov chain probabilistic modelling to describe how the decoded prefrontal attentional spotlight explores space in two different versions of a cued target detection task: a first version (the one used up to now), in which the cue could orient attention to one of the four possible quadrants (18 sessions), and a second version in which the cue oriented attention to only two possible quadrants, placed along the diagonal one with respect to the other (16 sessions). For both monkeys, oscillations in the prefrontal attention information did not depend on the task being performed by the animals (fig. 10A). In contrast, how the rhythmic attentional signal was deployed onto space as inferred from the decoded attentional spotlight was drastically different between the two tasks. This is captured by the Markov chain probabilistic modelling of the probability of the spotlight to stay in the cued quadrant when already there, or to shift to one of the uncued quadrants (fig. 10B, see methods). Indeed, while during the two types of task configurations, the probability that the decoded attentional spotlight remained at the cued location was highest (probabilities of 0.55 and 0.47 respectively), the pattern of probability of attention transitioning from the cued location to one of the uncued quadrants was very distinct. Specifically, during the four position task, virtually no transitions between the cued quadrant and the diagonally opposite quadrant can be observed (fig. 10B, grey, probability of transition of 0; for comparison, probability of transition from cued location to position 2: 22%; to position 3: 23%). This is exemplified in figure 10C, which represents the decoded attention spotlight trajectory during the cue to target interval in a representative trial of a four position task. In contrast, during, the two position task, transitions between the cued quadrant and the diagonally opposite quadrant, the

second most relevant spatial location in the task, become dominant with respect to the other two uncued quadrants (fig. 10B, red, probability of transition of 22%; for comparison, probability of transition from cued location to position 2: 16%; to position 3: 15%). This is exemplified in fig. 10D, which represents the decoded attention spotlight trajectory during the cue to target interval in a representative trial of a two position task.

Finally, we provide evidence to the effect that the prefrontal attentional spotlight explores space at an alpha that remains stable within trials and across tasks. In addition, we show that how this decoded spotlight explores space depends on both within-trial and across-task task contingencies.

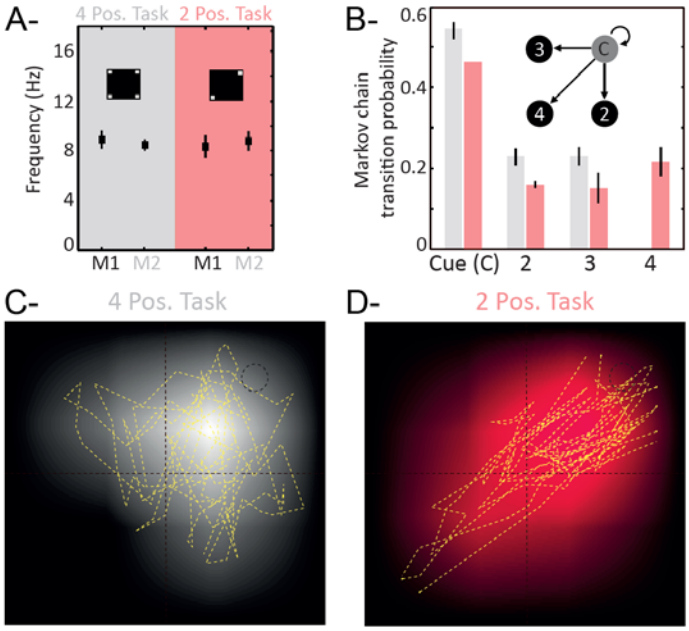

Figure 10: ***How prefrontal attentional spotlight rhythmically explores space depends on task context.*** (A) 7-12 Hz oscillation peak in prefrontal attention information in a four (gray) or two (red) cued positions target detection task, for monkeys M1 (black) and M2 (gray), over all sessions. (B) Markov chain probability of the attentional spotlight to stay at the cued quadrant (C), to transition to the same hemifield uncued quadrant (2), to transition to the opposite hemifield uncued quadrant (3) or to transition to the diagonally opposite uncued quadrant (4), in the four (gray) or two (red) cued positions target detection task. (C) Single trial example of prefrontal attentional spotlight exploration (10 ms resolution) during the cue to target interval, in a four cued positions target detection task,

477 superposed onto the Markov chain probability map for this cued condition. Cue was presented in the
478 upper right quadrant. (D) Same as in (C) for a two cued positions target detection task.

479

480 ## Discussion

481 Overall, we show that the attentional spotlight, decoded from cortical prefrontal FEF activities at
482 high temporal resolution, explores space rhythmically. This rhythmic exploration takes place in the 6-12
483 Hz (alpha) frequency range, independently of the ongoing task. Importantly, these oscillations of the
484 attentional locus determine both neuronal sensory processing, defining how much information is
485 available in the PFC about incoming sensory stimuli, and perception, determining whether these
486 incoming sensory stimuli are prone to elicit an overt behavioral response or not. From the spatial
487 perspective, this exploration of space corresponds to attentional saccades. These attentional saccades
488 explore space biased by both within-trial and across-task contingencies, implementing an alternation
489 between exploration and exploitation cycles.

490 ***The prefrontal attentional spotlight explores space rhythmically***

491 Converging behavioral evidence indicates that attention and perception are not anchored at a
492 specific location in space, but rather exhibit a temporal alpha rhythmicity (*42*). This rhythmic sampling of
493 space is phase-reset and entrained by external events of interest. It can also be observed spontaneously
494 (*43*), and is proposed to organize the tracking of task-relevant spatial locations by attention in time (*2*, *4*,
495 *40–42*, *44–46*). It has been proposed that, when prior information is available, such a rhythmic sampling
496 of information is more efficient than a continuous sampling of space (*47*). These observations have led to
497 reconsider the model of a continuously active attention spotlight in favor of a rhythmic sampling of
498 attention at relevant spatial locations, including during sustained attention states (*2*, *42*).

499 Our present findings reconcile these two models, describing a dynamic prefrontal attentional
500 spotlight that continuously explores space at a specific rhythm. This rhythmic exploration shares major
501 characteristics with previous behavioral reports on attentional rhythms: (1) these oscillations are
502 ongoing and can be identified independently of the intervening task events, (2) they are reset by
503 relevant external events such as spatial cues and (3) they occur in a well-defined functional alpha
504 frequency range. However, even if attentional exploration targets task relevant locations, as reflected by
505 the rhythmic enhancement of neuronal sensory processing and behavioral performance at the cued
506 target location, exploration is not restricted to these locations. Rather, space exploration by the

attention spotlight extends to un-cued *a priori* task irrelevant spatial locations, as reflected by the rhythmic enhancement of neuronal sensory processing and behavioral overt report at un-cued unpredictable distractor locations.

The phase between the attentional spotlight ongoing oscillations and a given stimulus presentation accounts from 10% (in the case of the target) to 30% (in the case of distractors) of the accuracy with which prefrontal neuronal populations encode the location of this stimulus. In other terms these oscillations – i.e. where the attention spotlight falls in space- critically impact the sensory processing of incoming stimuli. Tracing down this effect all throughout the visual system would be extremely relevant. Neuronal responses to low-salience task-relevant stimuli has been shown to arise earlier in the PFC than in the parietal cortex (Ibos et al. 2013). As a result, one predicts that this dependence of sensory processing onto attention spotlight oscillations will be found at all stages of the visual system. However, phase relationships between local neuronal and stimulus presentation is expected to vary, reflecting a top-down cascade of influences, in agreement with the role of the FEF in attentional control (*16*, *24–26*, *48*, *49*) .

These oscillations also determine overt behavioral perceptual outcome, accounting from 10% (in the case of false alarm production) to 30% (in the case of correct target identifications) of stimulus detection. This is globally higher than the range of reported oscillatory changes in behavioral hit rates (*1*, *2*, *5*), highlighting the high predictive power of these neuronal population oscillations.

Overall, this suggests the existence of perceptual cycles (*42*) that organize as a rhythmic alternation between exploitation and exploration states of space sampling by attention.

Exploring versus exploiting space by attention

Two models have been proposed to account for the spatial deployment of attention (*50–54*) a parallel processing model, driven by bottom-up information, dominating when visual search is easy; and a serial processing model, driven by top-down mechanisms, dominating in difficult visual search (*40*). In the context of this latter model, it has been hypothesized that the brain controls an attentional spotlight that scans space for relevant sensory information. In a previous study (*7*), we assessed, based on the (x,y) decoding of the neuronal population activity of the FEF, the tracking of this attentional spotlight in time (*7*). Here, we show that this prefrontal attentional spotlight explores space serially both at relevant (cued) and irrelevant (un-cued) locations, alternating between the exploitation and the exploration of the visual scene (*42*). The activity of the parietal (*19*) and prefrontal (*55*) cortical regions has been shown

to change drastically between exploitation and exploration behavior. In particular, exploration is associated with faster though less accurate oculomotor behavior (*19*) and a disruption of prefrontal control signals (*55*). This is proposed to facilitate the processing of unexpected external events (*19*), the expression of novel behavior and learning through trial and error (*55*).

Our observations strongly indicate that exploration and exploitation dynamically alternate within trials. This alternation of exploration and exploitation of space by the attentional spotlight thus appears to optimize subject's access to incoming information from the environment by a continuous exploration strategy, very much like is described for overt exploration behaviors such as saccadic eye movements, whisking or sniffing (*1*, *56*, *57*). This covert exploration of the environment by attention however takes place at a slightly higher frequency than the typical theta exploration frequency described for overt exploration. This is probably due to energetic and inertial considerations in controlling the remote effector during overt exploration (e.g. eye, whisker or nose muscles). Interestingly, the rhythm at which this prefrontal exploration/exploitation alternation takes place coincides with the rhythm at which attention behaviorally explores the different part of a given object (*2*). Overall, this leads us to postulate the existence of attentional saccades that can either be directed towards specific items for exploitation purposes, or deployed onto the entire visual field for exploration purposes.

Continuous attentional sampling and attentional saccades

Covert exploration of space by attention is more energy efficient than overt exploration by the eyes and the former serves to inform and guide the latter. In an initial "premotor theory of attention", these two processes, namely attentional selection and saccadic eye movements, have been suggested to rely on identical cortical mechanisms. This theory hypothesizes that attentional displacements or saccades of the mind, mirror saccades of the eyes except for the recruitment of the extra-ocular muscles (*58*). Since then, several studies have contributed to a functional dissociation between these two processes (*59–64*), and rhythmic attentional sampling has been shown to be independent from microsaccade generation (*5*, *65*, *66*). Our observations support a continuous exploration of space by the prefrontal attentional spotlight organized thanks to a rhythmic re-orientation of the attentional spotlight taking place at an alpha rhythm. This framework leads to an interesting set of experimental predictions. For instance, attentional capture and distractibility by an intervening distracting item is expected to coincide with an ongoing attentional re-orienting towards this item (*7*). Likewise, inhibition of return (*67–70*), is expected to reflect as an under-exploration of previously visited locations with respect to unexplored locations. This covert saccade-like exploration is proposed to be an intrinsic property of

attention, taking place irrespectively of the ongoing behavior and building onto a rhythmic alpha clock. Its spatial pattern, that is to say the portion of space that is being explored by these attentional saccades, as well as the frequency at which task-relevant items are explored are however expected to be under top-down control.

Top-down control

Numerous studies indicate that the PFC and specifically the FEF play a crucial role in attention orientation and attention control (*16*, *24–26*, *34*, *48*, *49*). As a result, one expects that the exploration of space by the prefrontal attention spotlight be strongly biased by top-down voluntary control. Confirming this prediction, we show that task goals significantly affect attentional space exploration strategy. Specifically, we observe that, the locations where the prefrontal attentional spotlight explores space are modulated both 1) within trials, by the expected position of the target after cue presentation, and 2) across tasks, by the general expectations about sensory events. In other words, relevant task items are more explored than irrelevant locations, where relevance concatenates information relative to the ongoing trial and task design. This is in agreement with prior behavioral observations reporting that the attentional sampling rate observed at the behavioral level decreases as the number of task relevant items increases (*3*, *41*). Overall, this indicates that the rhythmic exploration of space by attention, is an intrinsic, default-mode state of attention, that can be spatially modulated by task context and internal expectations. A strong prediction is that this rhythmicity in attentional spatial processing will directly impact attention selection processes in lower level cortical areas, through long-range feedback processes (*65*), possibly mediated by NMDA receptors (*71*).

Conclusion

Overall, our work describes for the first time the spatial and temporal properties of the population prefrontal attention spotlight. Specifically, it demonstrates a continuous exploration of space, that is mediated by attentional saccades that unfold at an alpha 7-12 Hz rhythm and that accounts for both neuronal sensory processing reliability and overt behavioral variability. Importantly, it bridges the gap between behavioral evidence of attentional rhythmic space sampling and local field attention related oscillatory mechanisms (*5*, *6*, *42*), revealing the neuronal population dynamics associated with rhythmic attentional sampling.

Acknowledgments

S.B.H was supported by ERC Brain3.0 #681978, ANR-11-BSV4-0011 & ANR-14-ASTR-0011-01, LABEX CORTEX funding (ANR-11-LABX-0042) from the Université de Lyon, within the program Investissements d'Avenir (ANR-11-IDEX-0007) operated by the French National Research Agency (ANR). R.V. was funded by ERC P-CYCLES #614244. We thank research engineer Serge Pinède for technical support and Jean-Luc Charieau and Fidji Francioly for animal care. All procedures were approved by the local animal care committee (C2EA42-13-02-0401-01) and the Ministry of research, in compliance with the European Community Council, Directive 2010/63/UE on Animal Care.

References

1. A. N. Landau, P. Fries, Attention samples stimuli rhythmically. *Curr. Biol.* **22**, 1000–1004 (2012).
2. I. C. Fiebelkorn, Y. B. Saalmann, S. Kastner, Rhythmic sampling within and between objects despite sustained attention at a cued location. *Curr. Biol.* **23**, 2553–2558 (2013).
3. R. VanRullen, Visual attention: a rhythmic process? *Curr. Biol.* **23**, R1110-1112 (2013).
4. L. Dugué, M. Roberts, M. Carrasco, Attention Reorients Periodically. *Curr Biol.* **26**, 1595–1601 (2016).
5. I. C. Fiebelkorn, M. A. Pinsk, S. Kastner, A Dynamic Interplay within the Frontoparietal Network Underlies Rhythmic Spatial Attention. *Neuron*. **99**, 842-853.e8 (2018).
6. R. F. Helfrich *et al.*, Neural Mechanisms of Sustained Attention Are Rhythmic. *Neuron*. **99**, 854-865.e5 (2018).
7. E. Astrand, C. Wardak, P. Baraduc, S. Ben Hamed, Direct Two-Dimensional Access to the Spatial Location of Covert Attention in Macaque Prefrontal Cortex. *Curr. Biol.* **26**, 1699–1704 (2016).
8. M. Posner, Orienting of Attention. *The Quarterly journal of experimental psychology*. **32**, 3–25 (1980).
9. M. Albares *et al.*, Attention to baseline: does orienting visuospatial attention really facilitate target detection? *J. Neurophysiol.* **106**, 809–816 (2011).
10. G. Ibos, J.-R. Duhamel, S. Ben Hamed, The spatial and temporal deployment of voluntary attention across the visual field. *PLoS ONE*. **4**, e6716 (2009).
11. H. S. Bashinski, V. R. Bacharach, Enhancement of perceptual sensitivity as the result of selectively attending to spatial locations. *Perception & Psychophysics*. **28**, 241–248 (1980).
12. Y. Yeshurun, M. Carrasco, Attention improves or impairs visual performance by enhancing spatial resolution. *Nature*. **396**, 72–75 (1998).
13. K. Anton-Erxleben, M. Carrasco, Attentional enhancement of spatial resolution: linking behavioural and neurophysiological evidence. *Nat. Rev. Neurosci.* **14**, 188–200 (2013).
14. C. Wardak, S. Denève, S. Ben Hamed, Focused visual attention distorts distance perception away from the attentional locus. *Neuropsychologia*. **49**, 535–545 (2011).
15. K. M. Armstrong, M. H. Chang, T. Moore, Selection and maintenance of spatial information by frontal eye field neurons. *J. Neurosci.* **29**, 15621–15629 (2009).
16. G. Ibos, J.-R. Duhamel, S. Ben Hamed, A functional hierarchy within the parietofrontal network in stimulus selection and attention control. *J. Neurosci.* **33**, 8359–8369 (2013).
17. C. J. McAdams, J. H. Maunsell, Effects of attention on orientation-tuning functions of single neurons in macaque cortical area V4. *J. Neurosci.* **19**, 431–441 (1999).
18. J. Lee, T. Williford, J. H. R. Maunsell, Spatial attention and the latency of neuronal responses in macaque area V4. *J. Neurosci.* **27**, 9632–9637 (2007).
19. S. Ben Hamed, J.-R. Duhamel, F. Bremmer, W. Graf, Visual receptive field modulation in the lateral intraparietal area during attentive fixation and free gaze. *Cereb. Cortex*. **12**, 234–245 (2002).
20. T. Womelsdorf, K. Anton-Erxleben, F. Pieper, S. Treue, Dynamic shifts of visual receptive fields in cortical area MT by spatial attention. *Nat. Neurosci.* **9**, 1156–1160 (2006).
21. M. R. Cohen, J. H. R. Maunsell, Attention improves performance primarily by reducing interneuronal correlations. *Nature Neuroscience*. **12**, 1594–1600 (2009).
22. C. W. Eriksen, J. D. St. James, Visual attention within and around the field of focal attention: A zoom lens model. *Perception & Psychophysics*. **40**, 225–240 (1986).
23. T. J. Buschman, E. K. Miller, Top-down versus bottom-up control of attention in the prefrontal and posterior parietal cortices. *Science*. **315**, 1860–1862 (2007).
24. C. Wardak, G. Ibos, J.-R. Duhamel, E. Olivier, Contribution of the monkey frontal eye field to covert visual attention. *J. Neurosci.* **26**, 4228–4235 (2006).

25. G. G. Gregoriou, S. J. Gotts, H. Zhou, R. Desimone, High-frequency, long-range coupling between prefrontal and visual cortex during attention. *Science*. **324**, 1207–1210 (2009).
26. L. B. Ekstrom, P. R. Roelfsema, J. T. Arsenault, G. Bonmassar, W. Vanduffel, Bottom-up dependent gating of frontal signals in early visual cortex. *Science*. **321**, 414–417 (2008).
27. E. Astrand, C. Wardak, S. B. Hamed, Neuronal population correlates of target selection and distractor filtering. *bioRxiv*, 422873 (2018).
28. C. A. Bosman, T. Womelsdorf, R. Desimone, P. Fries, A microsaccadic rhythm modulates gamma-band synchronization and behavior. *J. Neurosci.* **29**, 9471–9480 (2009).
29. H. Hogendoorn, Voluntary Saccadic Eye Movements Ride the Attentional Rhythm. *J Cogn Neurosci*. **28**, 1625–1635 (2016).
30. J. Otero-Millan, X. G. Troncoso, S. L. Macknik, I. Serrano-Pedraza, S. Martinez-Conde, *J Vis*, in press, doi:10.1167/8.14.21.
31. A. Wutz, E. Muschter, M. G. van Koningsbruggen, N. Weisz, D. Melcher, Temporal Integration Windows in Neural Processing and Perception Aligned to Saccadic Eye Movements. *Curr. Biol.* **26**, 1659–1668 (2016).
32. R. W. Berg, D. Kleinfeld, Rhythmic whisking by rat: retraction as well as protraction of the vibrissae is under active muscular control. *J. Neurophysiol.* **89**, 104–117 (2003).
33. E. E. Fanselow, M. A. Nicolelis, Behavioral modulation of tactile responses in the rat somatosensory system. *J. Neurosci.* **19**, 7603–7616 (1999).
34. E. Astrand, G. Ibos, J.-R. Duhamel, S. Ben Hamed, Differential dynamics of spatial attention, position, and color coding within the parietofrontal network. *J. Neurosci.* **35**, 3174–3189 (2015).
35. E. Astrand, C. Wardak, S. Ben Hamed, Selective visual attention to drive cognitive brain-machine interfaces: from concepts to neurofeedback and rehabilitation applications. *Front Syst Neurosci*. **8**, 144 (2014).
36. S. Farbod Kia, E. Åstrand, G. Ibos, S. Ben Hamed, Readout of the intrinsic and extrinsic properties of a stimulus from un-experienced neuronal activities: towards cognitive neuroprostheses. *J. Physiol. Paris*. **105**, 115–122 (2011).
37. A. Grinsted, J. C. Moore, S. Jevrejeva, Application of the cross wavelet transform and wavelet coherence to geophysical time series. *Nonlinear Processes in Geophysics*. **11**, 561–566 (2004).
38. S. Tremblay, G. Doucet, F. Pieper, A. Sachs, J. Martinez-Trujillo, Single-Trial Decoding of Visual Attention from Local Field Potentials in the Primate Lateral Prefrontal Cortex Is Frequency-Dependent. *J. Neurosci.* **35**, 9038–9049 (2015).
39. C. de Solages *et al.*, High-frequency organization and synchrony of activity in the purkinje cell layer of the cerebellum. *Neuron*. **58**, 775–788 (2008).
40. L. Dugué, D. McLelland, M. Lajous, R. VanRullen, Attention searches nonuniformly in space and in time. *Proc. Natl. Acad. Sci. U.S.A.* **112**, 15214–15219 (2015).
41. A. O. Holcombe, W.-Y. Chen, Splitting attention reduces temporal resolution from 7 Hz for tracking one object to <3 Hz when tracking three. *J Vis*. **13**, 12 (2013).
42. R. VanRullen, Attention Cycles. *Neuron*. **99**, 632–634 (2018).
43. N. A. Busch, J. Dubois, R. VanRullen, The phase of ongoing EEG oscillations predicts visual perception. *J. Neurosci.* **29**, 7869–7876 (2009).
44. L. Dugué, R. Vanrullen, The dynamics of attentional sampling during visual search revealed by Fourier analysis of periodic noise interference. *J Vis*. **14** (2014), doi:10.1167/14.2.11.
45. R. Gulbinaite, T. van Viegen, M. Wieling, M. X. Cohen, R. VanRullen, Individual Alpha Peak Frequency Predicts 10 Hz Flicker Effects on Selective Attention. *J. Neurosci.* **37**, 10173–10184 (2017).
46. R. VanRullen, T. Carlson, P. Cavanagh, The blinking spotlight of attention. *Proc. Natl. Acad. Sci. U.S.A.* **104**, 19204–19209 (2007).

47. C. E. Schroeder, P. Lakatos, Low-frequency neuronal oscillations as instruments of sensory selection. *Trends Neurosci.* **32**, 9–18 (2009).
48. T. Moore, M. Fallah, Microstimulation of the frontal eye field and its effects on covert spatial attention. *J. Neurophysiol.* **91**, 152–162 (2004).
49. G. Gregoriou, A. Rossi, L. G Ungerleider, R. Desimone, Lesions of prefrontal cortex reduce attentional modulation of neuronal responses and synchrony in V4. *Nature neuroscience*. **17** (2014).
50. A. M. Treisman, G. Gelade, A feature-integration theory of attention. *Cogn Psychol.* **12**, 97–136 (1980).
51. A. Zenon, S. Ben Hamed, J.-R. Duhamel, E. Olivier, Spatial and temporal dynamics of attentional guidance during inefficient visual search. *PLoS ONE*. **3**, e2219 (2008).
52. A. Zénon, S. Ben Hamed, J.-R. Duhamel, E. Olivier, *J Vis*, in press.
53. A. Zénon, S. Ben Hamed, J.-R. Duhamel, E. Olivier, Attentional guidance relies on a winner-take-all mechanism. *Vision Res.* **49**, 1522–1531 (2009).
54. J. M. Wolfe, Moving towards solutions to some enduring controversies in visual search. *Trends in Cognitive Sciences*. **7**, 70–76 (2003).
55. R. B. Ebitz, E. Albarran, T. Moore, Exploration Disrupts Choice-Predictive Signals and Alters Dynamics in Prefrontal Cortex. *Neuron*. **97**, 450-461.e9 (2018).
56. G. Buzsáki, *Rhythms of the Brain* (Oxford University Press, 2006).
57. J. Drewes, R. VanRullen, This is the rhythm of your eyes: the phase of ongoing electroencephalogram oscillations modulates saccadic reaction time. *J. Neurosci.* **31**, 4698–4708 (2011).
58. G. Rizzolatti, L. Riggio, I. Dascola, C. Umiltá, Reorienting attention across the horizontal and vertical meridians: evidence in favor of a premotor theory of attention. *Neuropsychologia*. **25**, 31–40 (1987).
59. M. Corbetta *et al.*, A common network of functional areas for attention and eye movements. *Neuron*. **21**, 761–773 (1998).
60. C. Wardak, E. Olivier, J.-R. Duhamel, A deficit in covert attention after parietal cortex inactivation in the monkey. *Neuron*. **42**, 501–508 (2004).
61. C.-H. Juan, S. M. Shorter-Jacobi, J. D. Schall, Dissociation of spatial attention and saccade preparation. *Proc. Natl. Acad. Sci. U.S.A.* **101**, 15541–15544 (2004).
62. K. G. Thompson, K. L. Biscoe, T. R. Sato, Neuronal basis of covert spatial attention in the frontal eye field. *J. Neurosci.* **25**, 9479–9487 (2005).
63. C. J. Bruce, M. E. Goldberg, Primate frontal eye fields. I. Single neurons discharging before saccades. *J. Neurophysiol.* **53**, 603–635 (1985).
64. Y. Kodaka, A. Mikami, K. Kubota, Neuronal activity in the frontal eye field of the monkey is modulated while attention is focused on to a stimulus in the peripheral visual field, irrespective of eye movement. *Neurosci. Res.* **28**, 291–298 (1997).
65. G. Spyropoulos, C. A. Bosman, P. Fries, A theta rhythm in macaque visual cortex and its attentional modulation. *Proc. Natl. Acad. Sci. U.S.A.* **115**, E5614–E5623 (2018).
66. A. N. Landau, H. M. Schreyer, S. van Pelt, P. Fries, Distributed Attention Is Implemented through Theta-Rhythmic Gamma Modulation. *Curr. Biol.* **25**, 2332–2337 (2015).
67. M. Posner, R. Rafal, L. S. Choatec, J. Vaughand, Inhibition of return: Neural Basis and Function. *Cognitive Neuropsychology*. **Vol. 2**, 211–228 (1985).
68. M. Posner, Y. Cohen, Components of visual orienting. *Attention and performance X: Control of language processes*. **32**, 531 (1984).
69. null Klein, Inhibition of return. *Trends Cogn. Sci. (Regul. Ed.)*. **4**, 138–147 (2000).

70. J. Lupiáñez, R. M. Klein, P. Bartolomeo, Inhibition of return: Twenty years after. *Cognitive Neuropsychology*. **23**, 1003–1014 (2006).
71. T. van Kerkoerle *et al.*, Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proc. Natl. Acad. Sci. U.S.A.* **111**, 14332–14341 (2014).

Supplementary figure

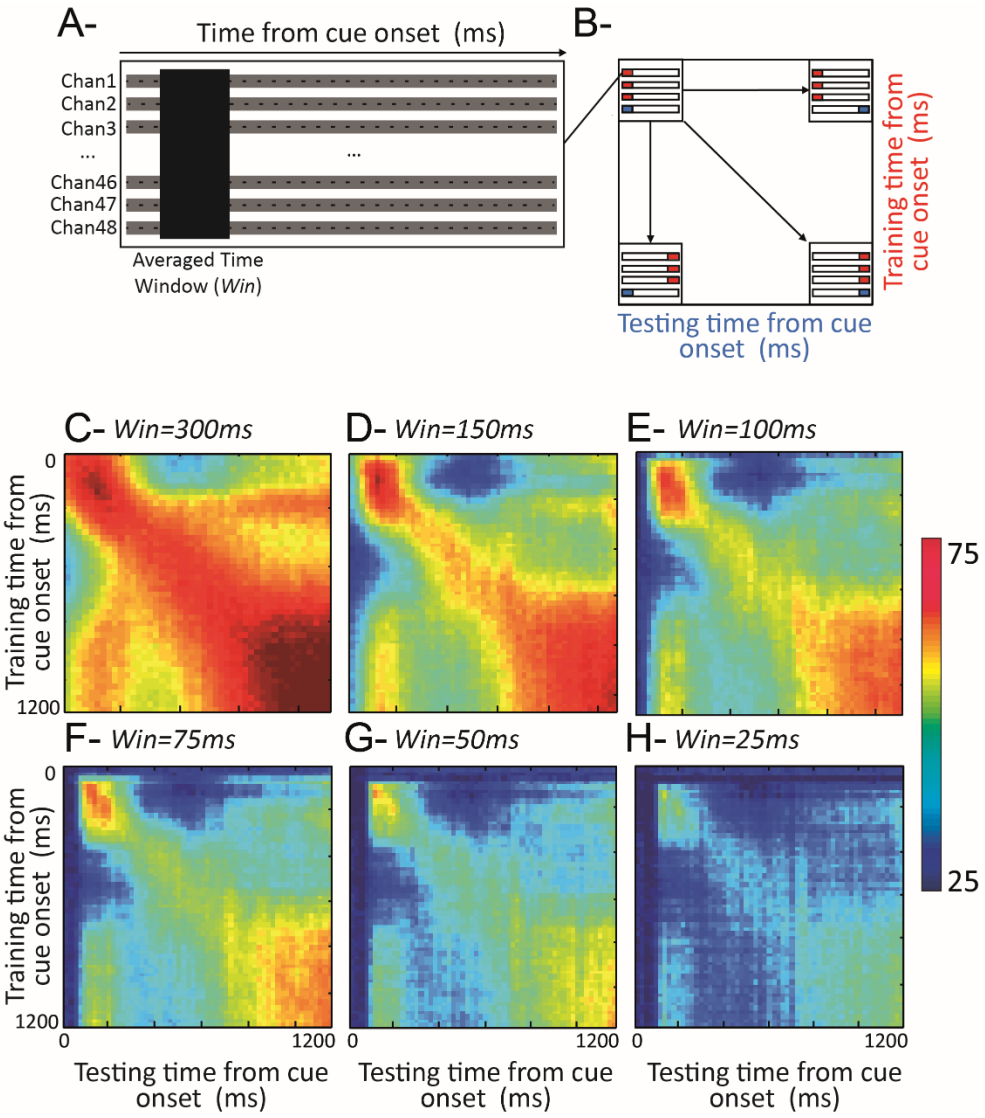

Figure S1: **Cross temporal decoding and the impact of averaging time windows.** (a) Data structure on a given trial: MUA activity is recorded onto 48 channels, in time (1 ms resolution), aligned with respect to the cue, and averaged over time windows of length *Win*. (b) Cross-temporal decoding matrices are obtained by training a decoder on activities from the 48-channels, collected at a given time *t,* averaged over ATW ms, on a subset of trials (random 70%) and testing this decoder on activities from the 48-channels, collected, averaged over ATW ms, at all possible times (resolution of 10ms), on the remaining 30% test trials. This procedure is repeated over and over by moving reference time *t* by 10 ms each time,

from 0 ms to 1200 ms from cue presentation. (c-h) Cross-temporal decoding matrices with different averaging time windows from 300 ms (c), to 150 ms (d), to 100 ms (e) , to 75 ms (f) , to 50 ms (g) , to 25 ms (h) averaging window. 50 ms averaging windows reveals oscillations in the decoding performance along the testing time dimension (x-axis). These oscillations can already be seen at *Win*=75 ms