Replay of rule-learning related neural patterns in the prefrontal cortex during sleep

A. Peyrache, M. Khamassi, K. Benchenane, S. I. Wiener, F. P. Battaglia

SUPPLEMENTARY DISCUSSION

Memory replay and network spikes.

The transient co-activations which make up most of replay resemble network spikes, transient activations of groups of cells, which have been described in mathematical models of recurrent neural networks, in which dynamics is governed by synaptic facilitation and depression [1]. These events are an intrinsic feature of network dynamics in several neural systems, e.g. organotypic cultures [2] and in neocortical slices [3]. From a theoretical viewpoint, network spikes may be triggered by small fluctuations in activity, activating a subset of cells encoded in the synaptic matrix, then terminating because of synaptic depression. Network spike firing may also, through synaptic facilitation, increase the likelihood of a successive activation of the same group of cells [1] in a self-sustained, repetitive process, favoring synaptic plasticity and a more permanent encoding of that cell assembly in the connectivity matrix. Thus, by themselves, network spikes do not denote learning or information processing. However, they provide a mechanism for the expression of learning in neural activities: patterns of activity may be tagged, for example by synaptic facilitation, or long-term synaptic plasticity, and be more likely to re-emerge later under the form of a network spike. Moreover, several groups of cells likely to give rise to network spikes may co-exist, and activate at different moments during sleep. Conceivably, the sleep activity of the mPFC could be largely composed of network spikes, each involving a different cell group. Of these groups, our technique can detect but a few among those that are related to the immediately preceding experience. Another interesting point is that different reactivating cell groups are unlikely to be active at the same time, as shown by the cross-correlograms for reactivation strengths relative to different PCs (Supplementary Fig. 10). This suggests that some sort of pattern separation mechanism (possibly through feedback inhibition), takes place specifically during sleep. This result does not contradict studies [4, 5] demonstrating the replay of neuronal ensemble activation in sequential order, spanning up to several hundreds of ms: our technique will search for the patterns of co-activation accounting for the largest fraction of the activity variance in the AWAKE epoch. With respect to sequences, two situations may arise. First the sequence is shorter than the bin size we used (100 ms), then it would be completely contained in one of our co-activation patterns. For longer sequences, or in any case sequences that are cut in two at the bin boundaries, nothing ensures that all parts of the sequence will be detected as co-activation patterns by our analysis, so that the sequential character of replay can be analyzed. In general this will not be the case, because here PCs are computed by AWAKE analysis without reference to precise, repetitive behavioral templates, as it is the case for analyses of sequential activation of ensembles. Thus, the present method and sequence-based methods highlight complementary aspects of sleep replay.

References

[1] Mongillo, G., Barak, O., and Tsodyks, M. *Science* **319**(5869), 1543–1546 March (2008).

[2] Eytan, D. and Marom, S. *J. Neurosci.* **26**(33), 8465–8476 August (2006).

[3] Cossart, R., Aronov, D., and Yuste, R. *Nature* **423**(6937), 283–288 May (2003).

[4] Ji, D. and Wilson, M. A. *Nature Neuroscience* **10**(1), 100–107 December (2006).

[5] Euston, D. R., Tatsuno, M., and Mcnaughton, B. L. *Science* **318**(5853), 1147–1150 November (2007).

[6] Buzsáki, G., Penttonen, M., Nádasdy, Z., and Bragin, A. *Proc Natl Acad Sci U S A* **93**(18), 9921–9925 September (1996).

[7] Henze, D. A., Borhegyi, Z., Csicsvari, J., Mamiya, A., Harris, K. D., and Buzsáki, G. *J Neurophysiol* **84**(1), 390–400 July (2000).

[8] Yokota, T. and Satoh, T. *Brain research bulletin* **54**(5), 575–584 March (2001).

[9] Fujisawa, S., Amarasingham, A., Harrison, M. T. T., and Buzsáki, G. *Nature neuroscience* May (2008).

SUPPLEMENTARY FIGURES

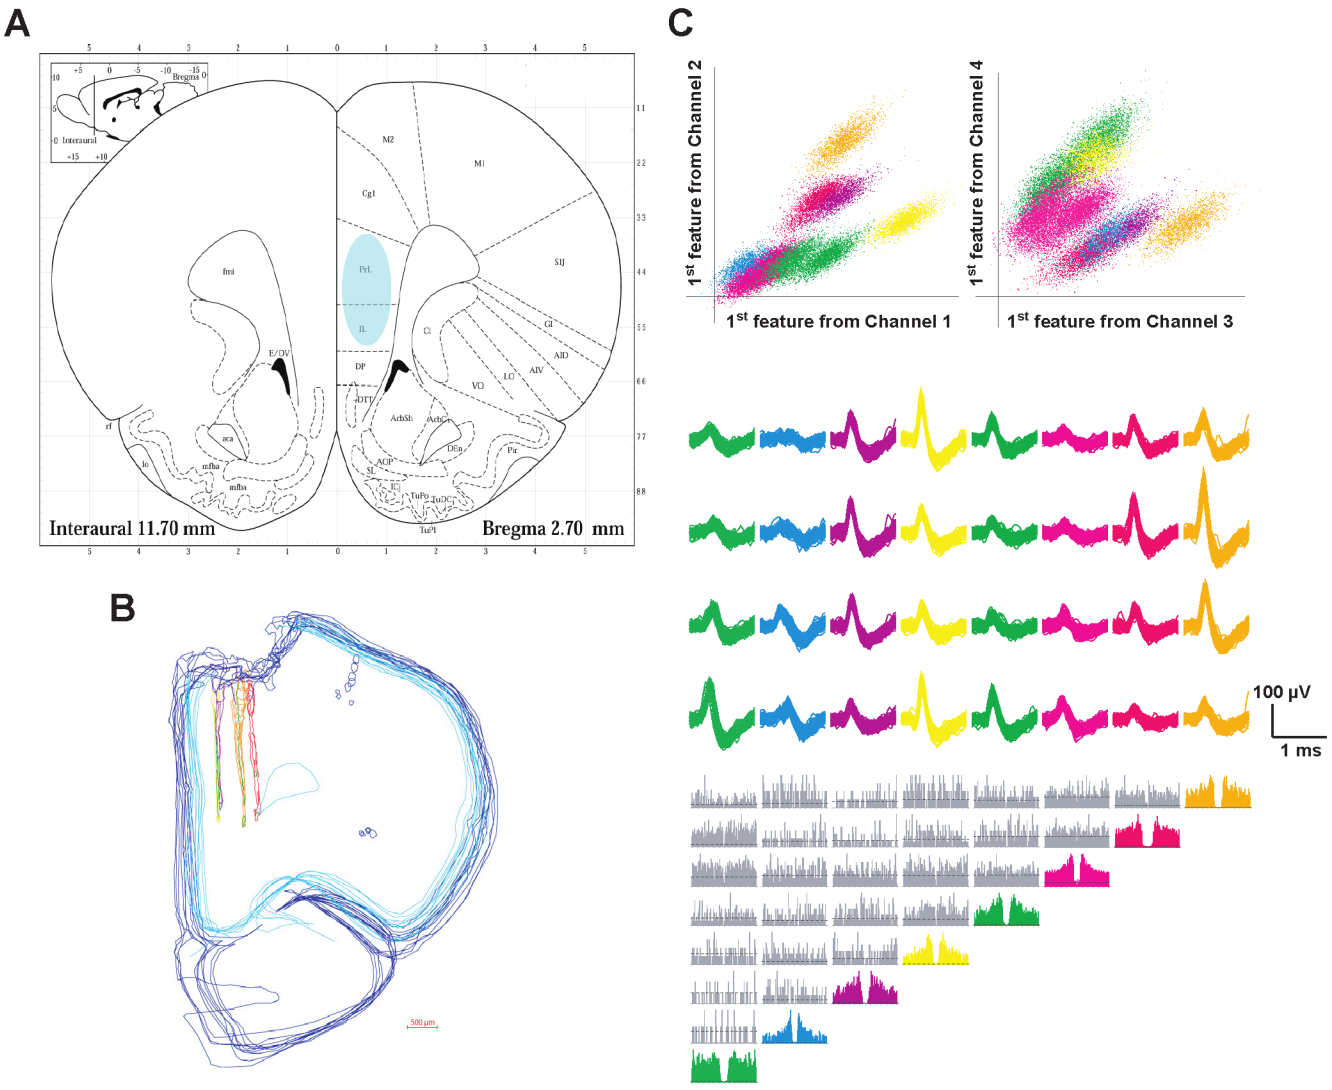

Supplementary Figure 1: Histology and spike sorting. A: recordings were made in the prelimbic and infralimbic area (blue circle) of rats (adapted from Paxinos and Watson 1998, "The rat brain in stereotaxic coordinates") B: reconstructed track of six tetrodes implanted in one rat. C. Example of 8 neurons recorded on the same tetrode (each associated with a different color). *Top* : each spike is represented in 4 of the 12 dimensions used for clustering of spikes (*left*: one feature from channel 1 versus one feature from channel 2; *right*: one feature from channel 3 versus one feature from channel 4, see METHODS), showing little overlap between these 4 dimensions among the clusters. *Middle* : Superimposed traces of spikes clustered accordingly to the above clustering. *Bottom*: Auto-correlograms (in colors) and cross-correlograms (grey bars) of cells' spike time showing clear refractory periods for each individual cell but not for cross-correlograms, hence showing little spike contamination within tetrodes (Displays from Klusters, see METHODS). While the second cluster (second column, cyan) displays slighty different waveforms, and different peak times, on the third and fourth channels, the auto- and cross-correlograms indicate good isolation. Furthermore such differences in waveforms have been demonstrated for electrodes situated at different distances from the soma (see for example [6]). This corresponds to the configuration here since tetrodes were implanted parallel to the medial bank, that is perpendicular to dendritic trees. However, such cases were rare in our datasets.

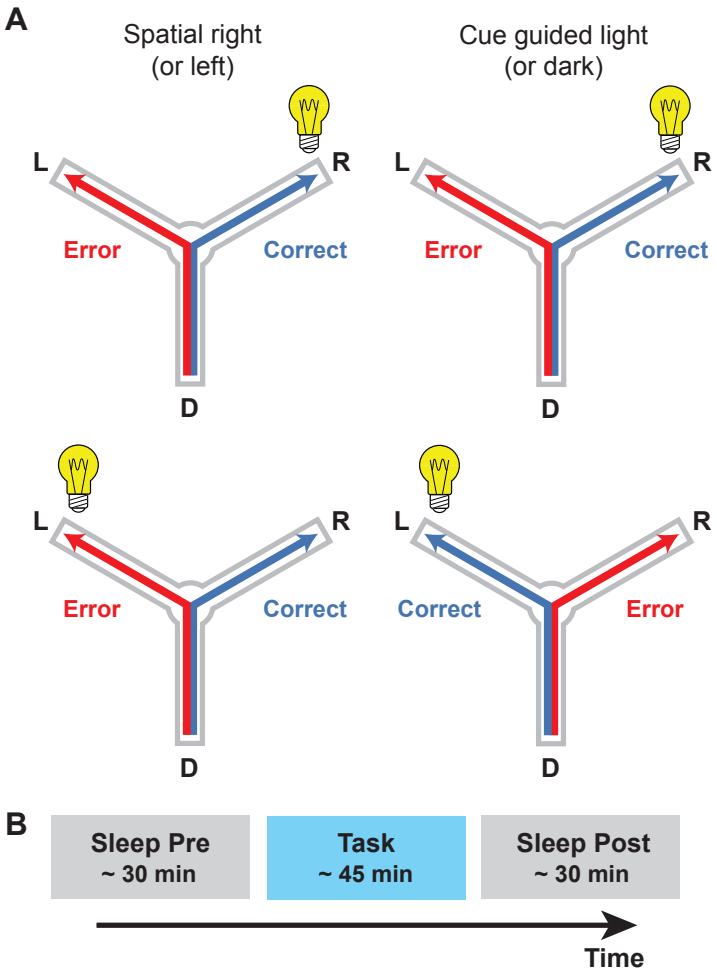

Supplementary Figure 2: Decision making task on the Y maze. A: Depiction of the 4 rules used for the behavioral task on the Y-maze (see METHODS). B: Timeline of a typical experimental session, including rest (gray) and task (cyan) epochs. D- Start arm; L Left goal arm; R-Right goal arm.

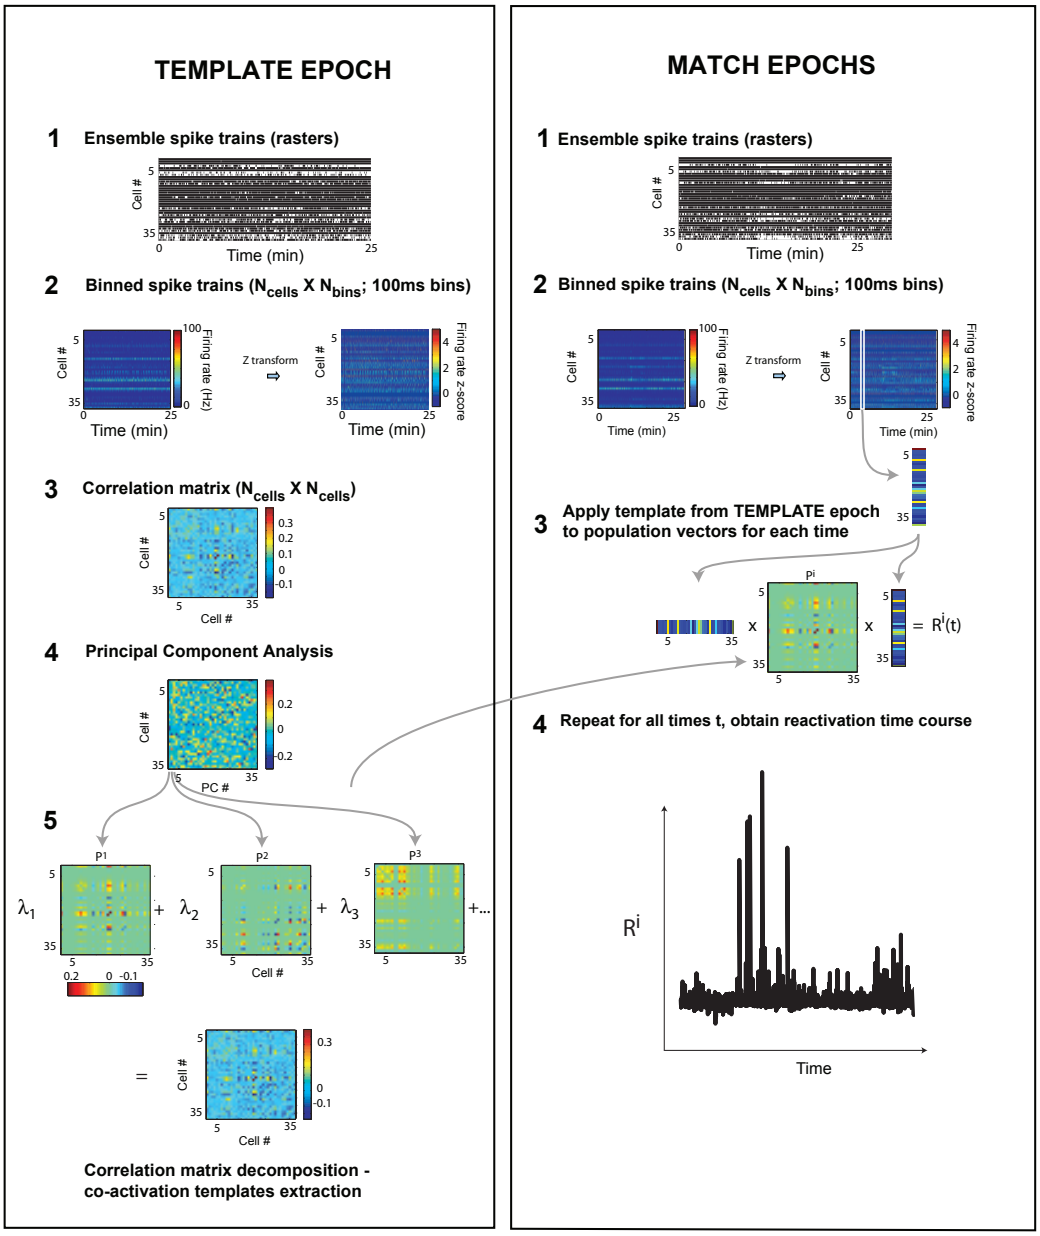

Supplementary Figure 3: The procedure used to compute reactivation strengths

TEMPLATE epoch: 1) Spike trains from ensembles of simultaneously recorded mPFC cells are obtained. 2) The spike trains are binned in 100 ms time windows (left) and then z-transformed, so that each binned spike train has zero mean and unit variance. 3) The correlation matrix of the z-transformed binned spike trains is computed (an $N_{cells} \times N_{cells}$ matrix). 4) The principal components of the matrix are computed: the matrix is diagonalized, and the eigenvectors are sorted according to their associated eigenvalues. 5) From each principal component, a projector operator, $P^1 \dots P^n$ is computed, as the external product of the component vector with itself. The sum of these projector operators, weighted by their respective eigenvalues, yields the correlation matrix. The projectors $P^1 \dots P^n$ are the templates that are matched to the activity in the rest epochs.

MATCH: Steps 1) and 2) as above. 3) Each template from the TEMPLATE epoch is separately applied on the MATCH activity: for each time t the population vector (binned spike count for each cell) is left and right multiplied on the projector matrix, and a value is obtained representing the degree of matching between the population vector and the TEMPLATE template. This is defined as the *reactivation strength* at time t. 4) The procedure is repeated for each time bin in the MATCH epoch, yielding a time course for the reactivation strength.

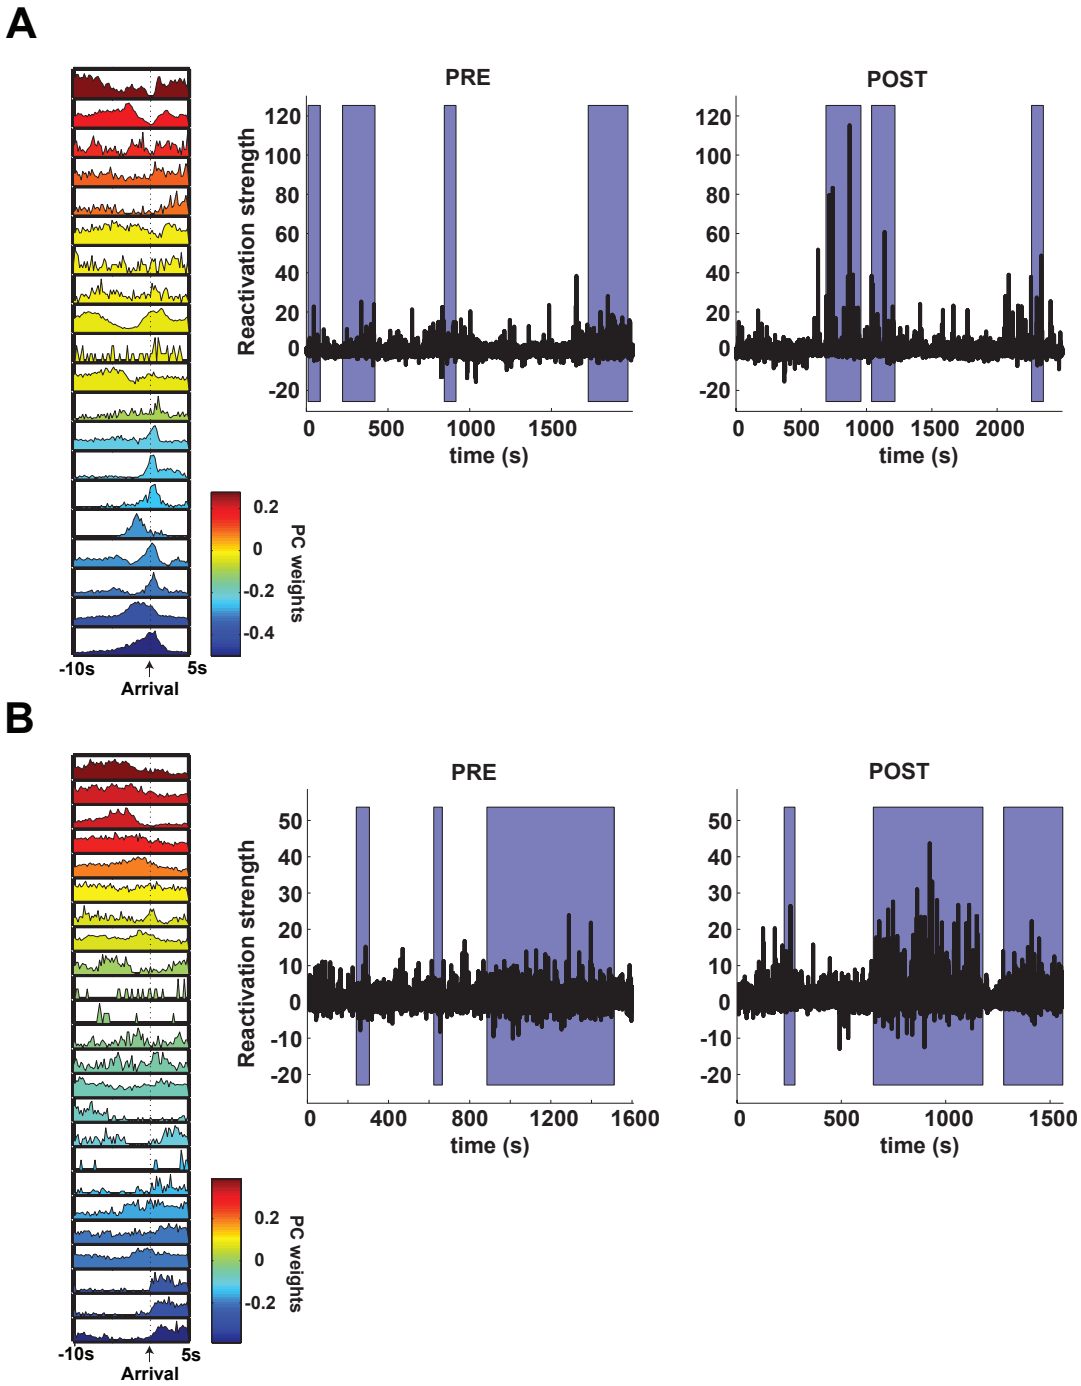

Supplementary Figure 4: A-B: Two more examples (from two different rats) of behavioral correlates of signal components and reactivation time course, comparable to what was presented in main text Fig. 1 and 2. (Left) Peri-event time histograms (PETH) of cells during the performance of the behavioral task (AWAKE epoch), synchronized with the time of arrival at the reward sites, presented in the same form as Fig. 1A: cells are presented in the order of their weight in the considered principal component. As in Fig. 1A, cells with similar eigenvalues have highly correlated PETHs. This suggests that the correlation matrix for 100 ms binned spike trains is largely dominated by correlation at longer time scales, which might be associated to behavioral effects. Note that, going from cells with positive to negative principal component weight, the task phase with increased activity does not necessarily have a sequential order in the task, that is, from trial start to trial end. (Right) Time course of the reactivation strengths during the PRE and POST rest epochs, as in Fig. 2A. Blue boxes delimit SWS epochs. As in the case of Fig. 2A, during POST (but not PRE) SWS, peaks in reactivation strength are present, corresponding to elevated reactivation.

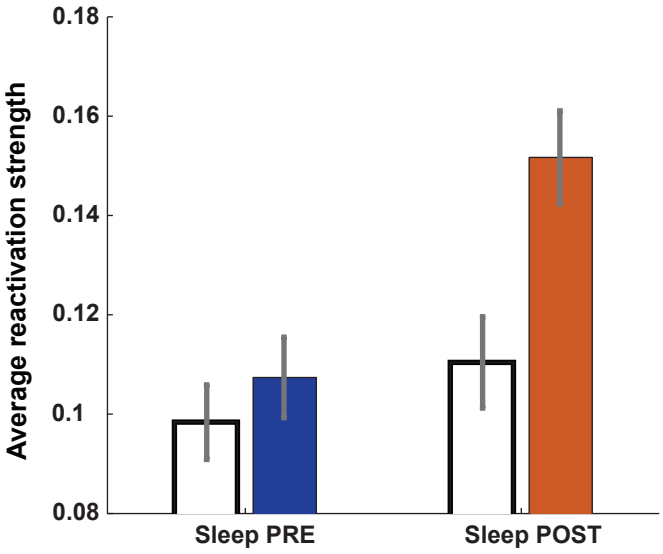

Supplementary Figure 5: Comparison of reactivation strength of signal components for SWS (colored bars) and non SWS epochs (white bars). A two-way ANOVA revealed significant effects of both SWS vs non-SWS ($p < 0.01$) and POST VS. PRE ($p < 0.01$). Bars indicate mean values ± SEM.

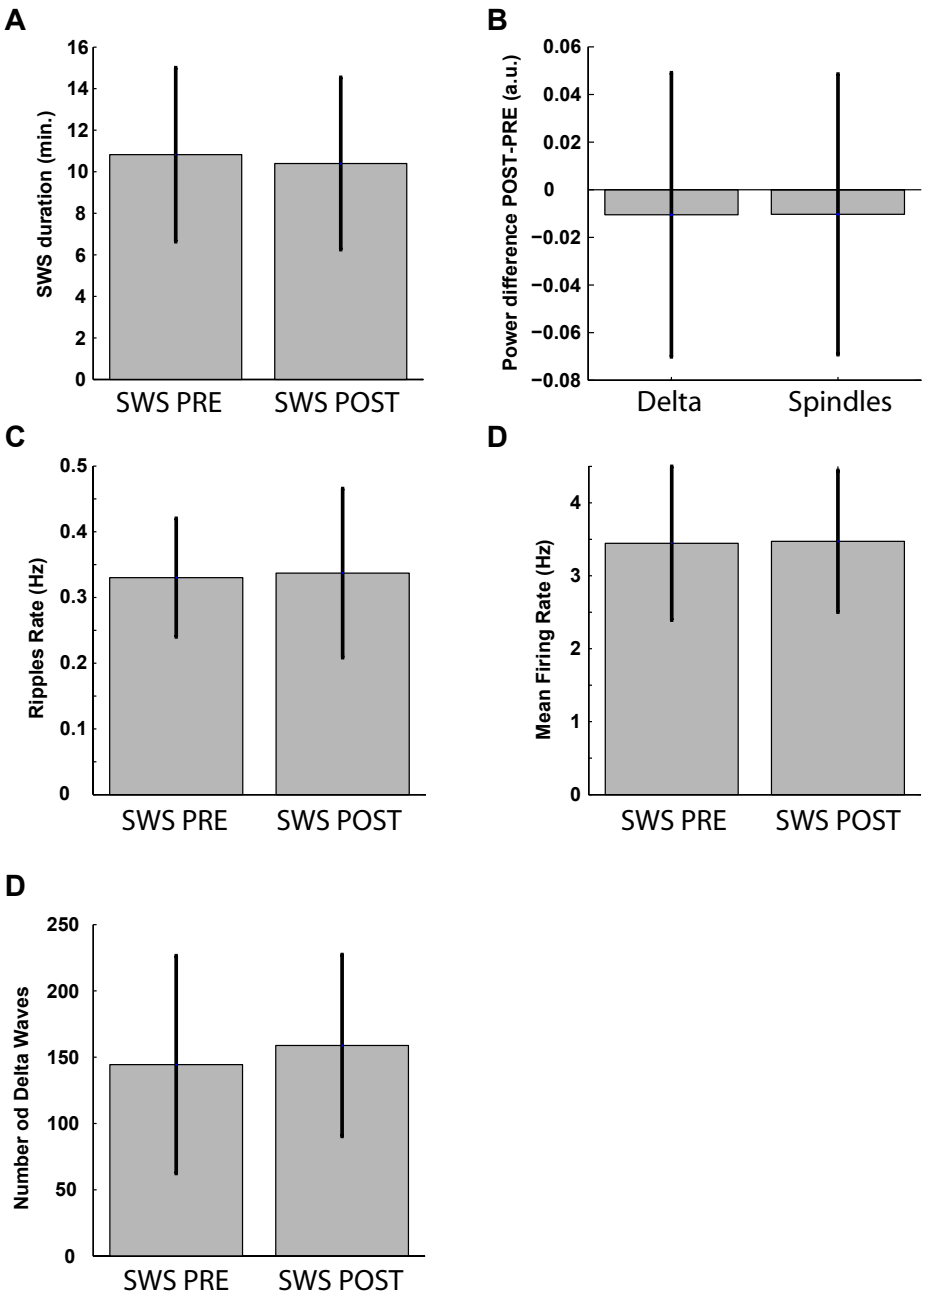

Supplementary Figure 6: Absence of significant differences between PRE and POST SWS in global measures of activity and oscillations (paired t-test, $p>0.05$). A. Average SWS duration; B. Averaged power difference in the delta band (1-4Hz) and spindles band (10-20 Hz) between POST and PRE SWS for each session. Only the difference is shown since the power scale is somewhat variable between sessions and thus poorly informative. C. Rate of ripples occurence in SWS. D. Average firing rates of cells. D. Number of detected Delta waves per session. Error bars display SD.

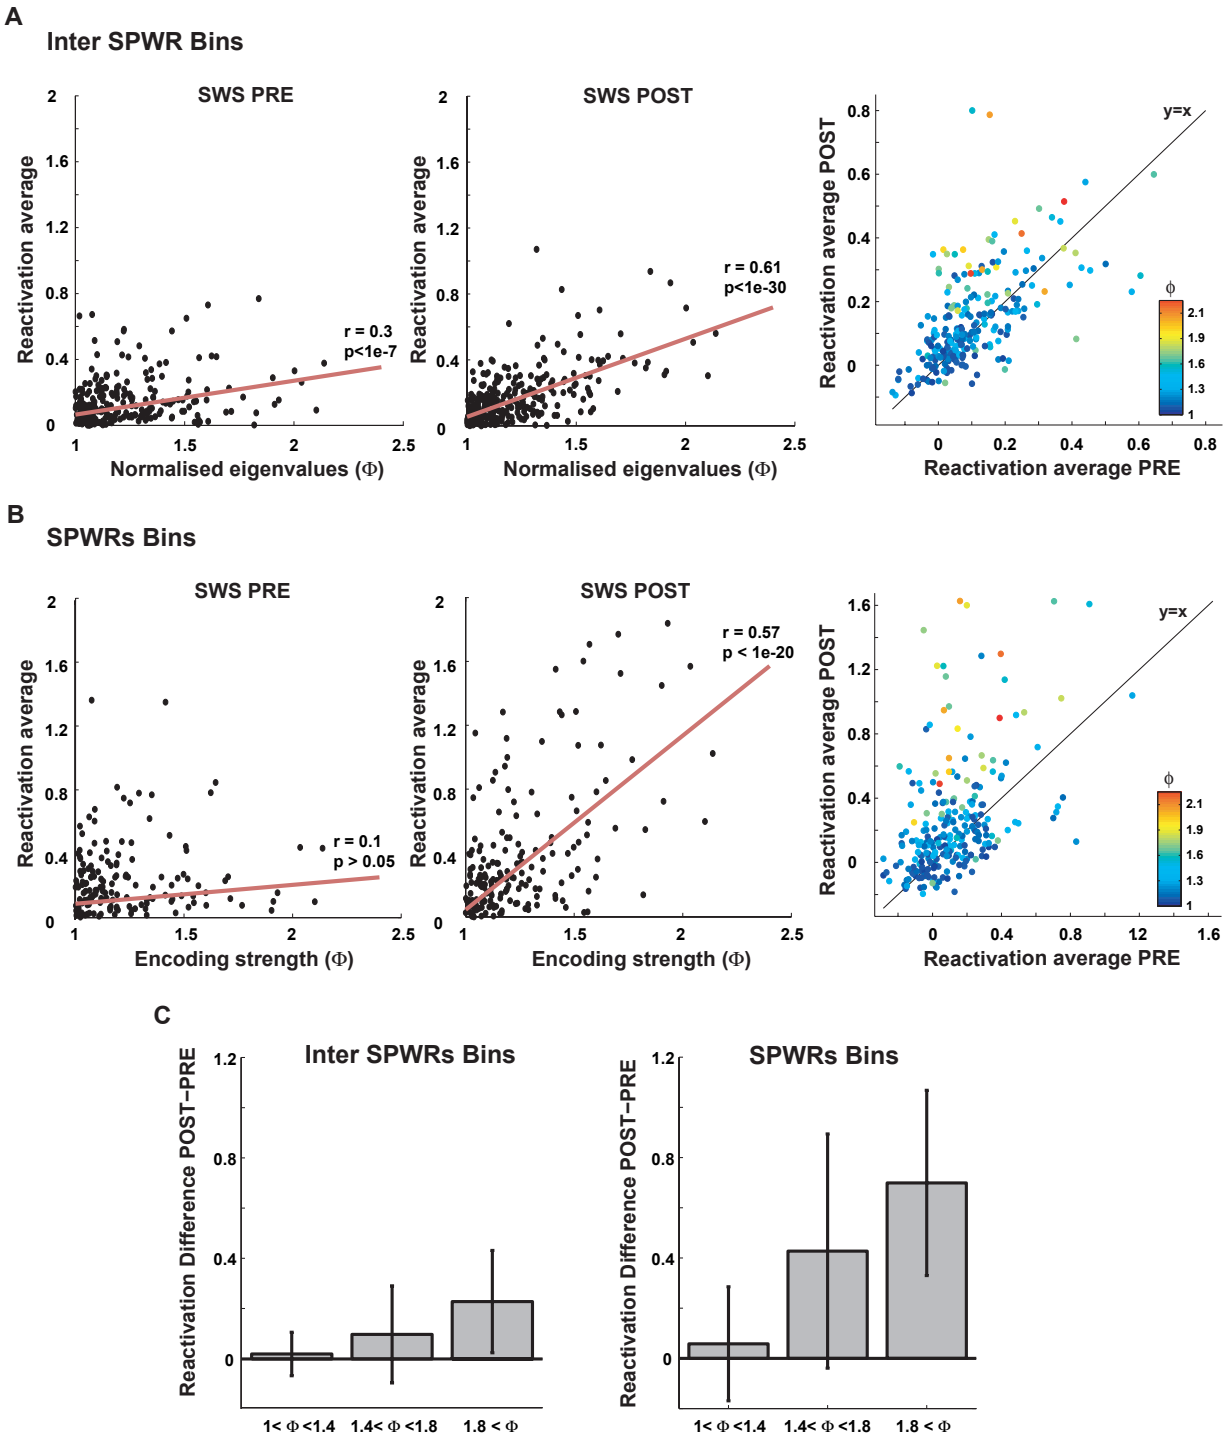

Supplementary Figure 7: Reactivation strength average increased with encoding strength. A. *Left-Middle*: Scatter plots of reactivation strength average for inter-SPWR bins as a function of normalized eigenvalues. POST SWS shows a stronger correlation. *Right*: Averaged reactivation strength during POST SWS vs. PRE SWS. Color scale: encoding strength. B: same plot but only bins coinciding with a hippocampal SPWR were considered. The correlation was significant only for POST (Pearson's correlation test, $p < 10^{-23}$). PCs with higher encoding strength tend to show the largest positive difference in replay between POST and PRE. C. Averaged reactivation strength for each group of PCs during inter-SPWRs bins (*left*) and SPWRs bins (*right*). Error bars display SD.

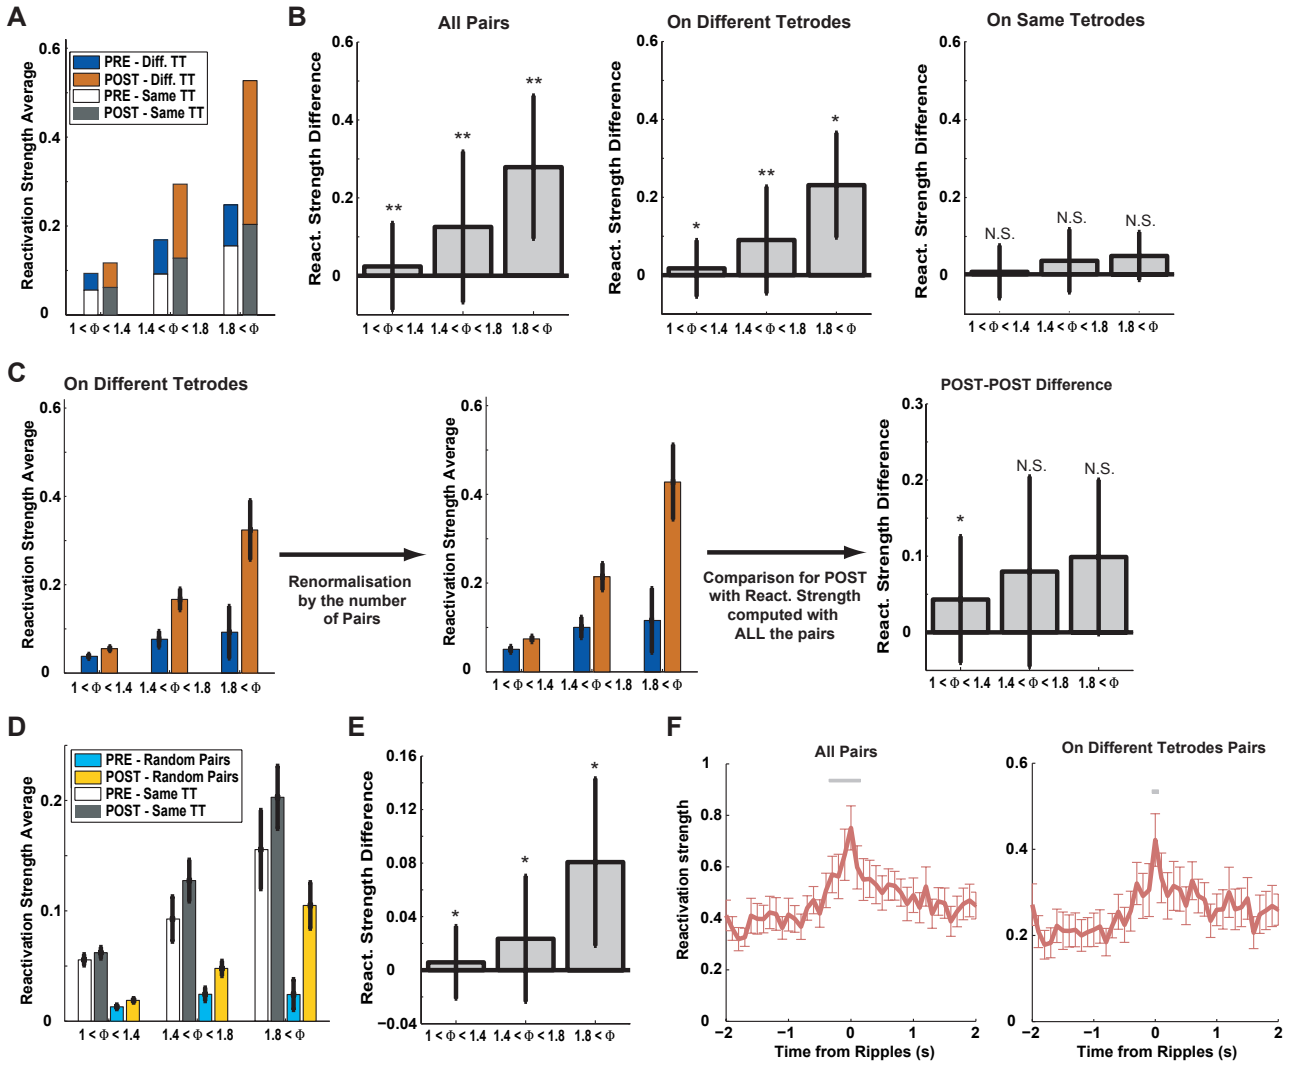

Supplementary Figure 8: Comparable results were obtained when pairs of cells recorded on the same tetrodes were excluded from the computations, controlling for possible artifacts due to erroneous spike sorting.

A. Bar plots of the epoch-wide average reactivation strengths, grouped by their normalized encoding strength $\Phi = \lambda/\lambda_{max}$, for PRE and POST. The reactivation strength averages for all pairs of neurons (same as in Fig. 2B) is decomposed in two different contributions: reactivation strength from on-same-tetrode pairs (below) and on-different-tetrode pairs (above). B. While the paired difference between POST and PRE epochs were significantly different when all cell pairs are considered (*left*), and for on-different-tetrode pairs (*center*), a significant difference was not observed for on-same-tetrode pairs (*right*). Thus, the difference we observe in pattern activation between PRE and POST is due to on-different-tetrode cell pairs only. C. Another issue is whether pattern activation strengths in the PRE and POST epoch are affected by the exclusion of on-same-tetrode cell pairs. To test this, reactivation strength can be computed from on-different tetrode cell pairs only (*left*, as in A), and then normalized to take into account the smaller number of terms in the sum (*center*, see below). The distribution of differences between the all pairs and the renormalized on-different-tetrode reactivation strength for POST epochs (*right*) show no significant difference for the highest encoding strength components ($\Phi > 1.4$). Only for the smallest encoding strength components $\Phi < 1.4$ were the subtracted reactivation strengths significantly (but slightly) smaller that the strengths computed with the full sum. D. The decomposition of reactivation strength shown in A, and the resulting comparison between on-same-tetrode and on-different-tetrode pairs are possibly affected by the larger number of on-different-tetrode pairs than on-same-tetrode pairs. To control for this, we considered on-same-tetrode pairs alone, and compared them, for each session, with 100 randomly chosen, matched-size groups of on same different tetrode pairs. For these two data sets, we computed average reactivation strengths for PRE and POST. Only for on-different-tetrode pairs (E) was there a

significant difference between POST and PRE, denoting replay. No difference was observed between PRE and POST for on-same-tetrode pairs. Notice that the absolute values of reactivation strengths were greater for on-same-tetrode cell pairs than on-different-tetrode pairs. This could in principle be due to spurious correlations introduced by spike sorting. However, cells recorded from the same tetrode are likely to be located less than $\sim 100\mu$m from each other [7], whereas cells on different tetrodes are at least 300μm (horizontal separation) away from each other, and usually much more, because of the differences in depth between tetrodes. Several studies show that the probability of a monosynaptic contact between two cells is much larger for cells closer to each other than 200μm than for more distant cells (see e.g. [8, 9]). Monosynaptic contacts would have the effect of imposing a strong, stationary correlation structure which would be detected in all three epochs, contributing to increased reactivation strengths values for both POST and PRE. Whether this result depends on spike sorting errors or on the intrinsic connectivity structure, it seems to have little or no bearing on the replay effects, since in no analysis is replay detected from on-same-tetrode cell pairs alone. F. The Event Triggered Average of reactivation strength (not renormalized) centered on hippocampal SPWRs, averaged over signal components with strong encoding strength ($\Phi > 1.4$) is shown for all cell pairs (*left*) and pairs of cells from different tetrodes only (*right*). Although the baseline was lower in the tetrodes exclusion condition, reactivation strength was significantly above baseline (t-test, $p<0.001$, indicated by grey bars) in the same time intervals. Error bars display SEM in C (*left* and *middle*), D and F, SD in B, C (*right*) and E.

Reactivation strengths were therefore computed by omitting pairs from the same tetrodes, leading to time series R_l^{noTT}. More precisely, the elements P_{ij}^l of the projector were replaced by 0s when cells i and j were on the same tetrode. Nevertheless, as the principal component is not renormalized in this operation, the overall reactivation strength was lowered. To compare obtained values with or without those pairs of neurons, R_l^{noTT} has to be renormalized by a factor F so that the quantities R_l and $F.R_l^{noTT}$ are comparable. As the number of pairs involved is decreased, the renormalisation factor F_{number} will be the ratio between the total number of pairs and the number of pairs on different tetrodes. If N_t is the total number of neurons, the total number of pairs is $(N^2-N)/2$. Similarly, n_i being the number of neurons on the ith tetrode, the number of on-same-tetrode pairs is $(\sum_i n_i^2 - n_i)/2$. Hence, the renormalisation factor is: $F = \frac{N^2-N}{(N^2-N)-\sum_i\left(n_i^2-n_i\right)}$

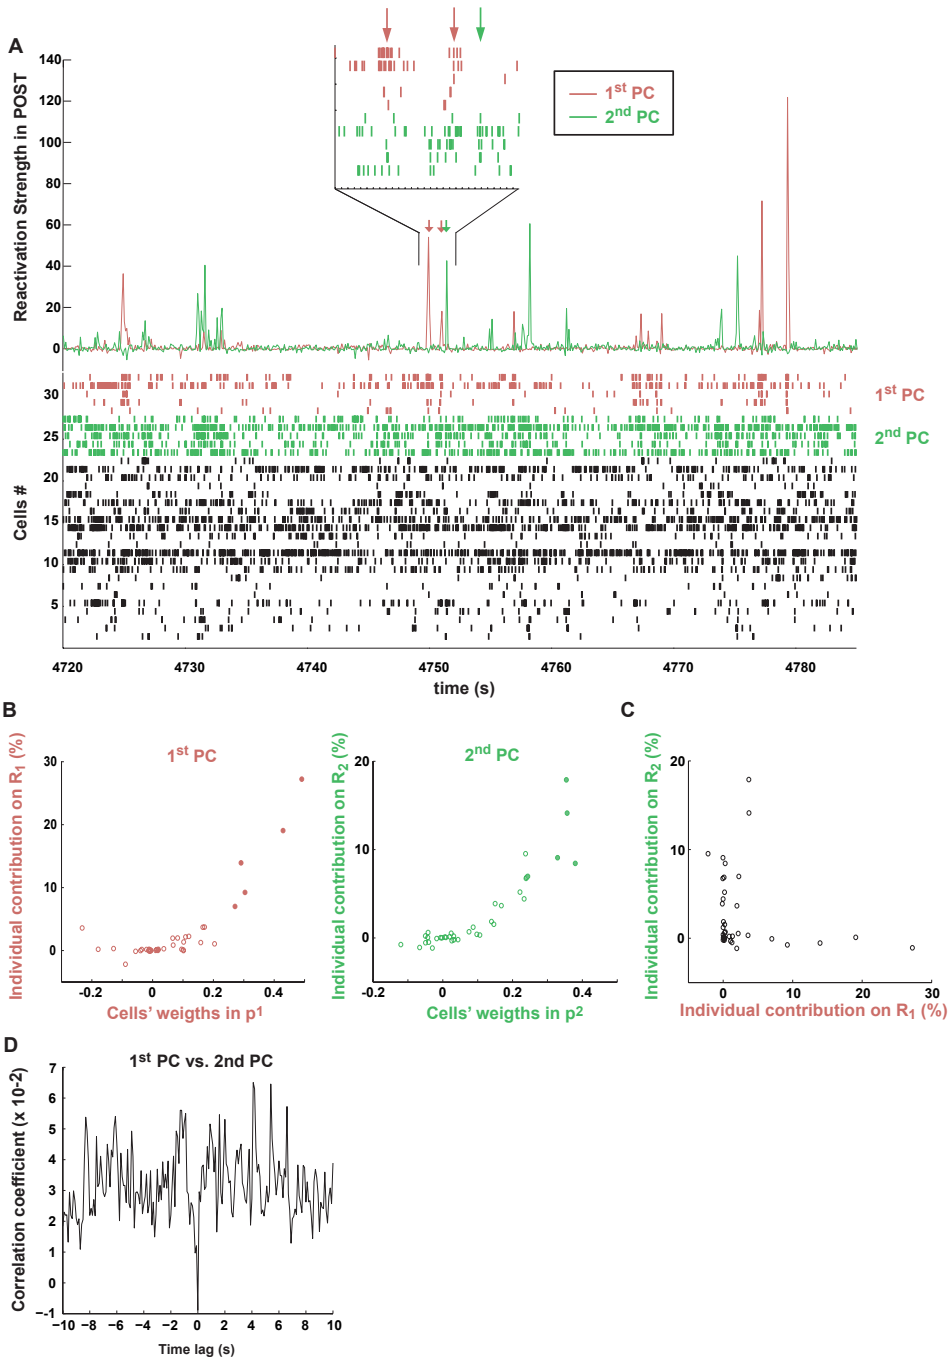

Supplementary Figure 9: Example of interactions between two different signal components

A. Time course of reactivation strength for two different principal components (*top*) and cells ensemble raster plot (*bottom*) during a POST SWS epoch. For each PC, spikes from the five cells with the highest weights in the PCs were sorted in the raster according to their weights in descending order. The other cells were plotted in no particular order. In correspondence with reactivation strength peaks, cells with the highest weights in the PC co-activate transiently, as shown in the expanded inset. B. Contribution of each individual cells to the overall reactivation strength during POST SWS (see below) as a function of the cell weight in the PC, following an approximately quadratic curve. The cells highlighted in the raster plot in A are represented by filled dots. C. Contribution of each cell to the reaction strength of the second PC in function of their contribution in the first. Note that two non-overlapping sets of cells

account have the strongest contribution in the two PCs D. Cross correlogram between the two reactivations during this same SWS epoch, showing a trough for zero lag, which implies that the two PCs are unlikely to reactivate at the same time.

Each neuron's spike train participates differently to any particular reactivation strength, likely depending on its associated weight in the principal component. To quantify the contribution of the kth cell, the reactivation strength R_l^{-k} is computed over $\forall t$, $Q_{kt} = 0$, or by removing the terms depending on Q_{kt}: $R_l^{-k}(t) = \sum_{\substack{i,j:i\neq j \\ i\neq k, j\neq k}} Q_{it}\mathbf{P}_{ij}^{(l)}Q_{jt}$ then the contribution is defined as $I_l^k = \frac{1}{2}\left(1 - \frac{\langle R_l^{-k}\rangle}{\langle R_l\rangle}\right)$ where $\langle . \rangle$ denotes the average over time and the normalisation factor $1/2$ has been derived from simple calculation so that $\sum_k I_l^k = 1$

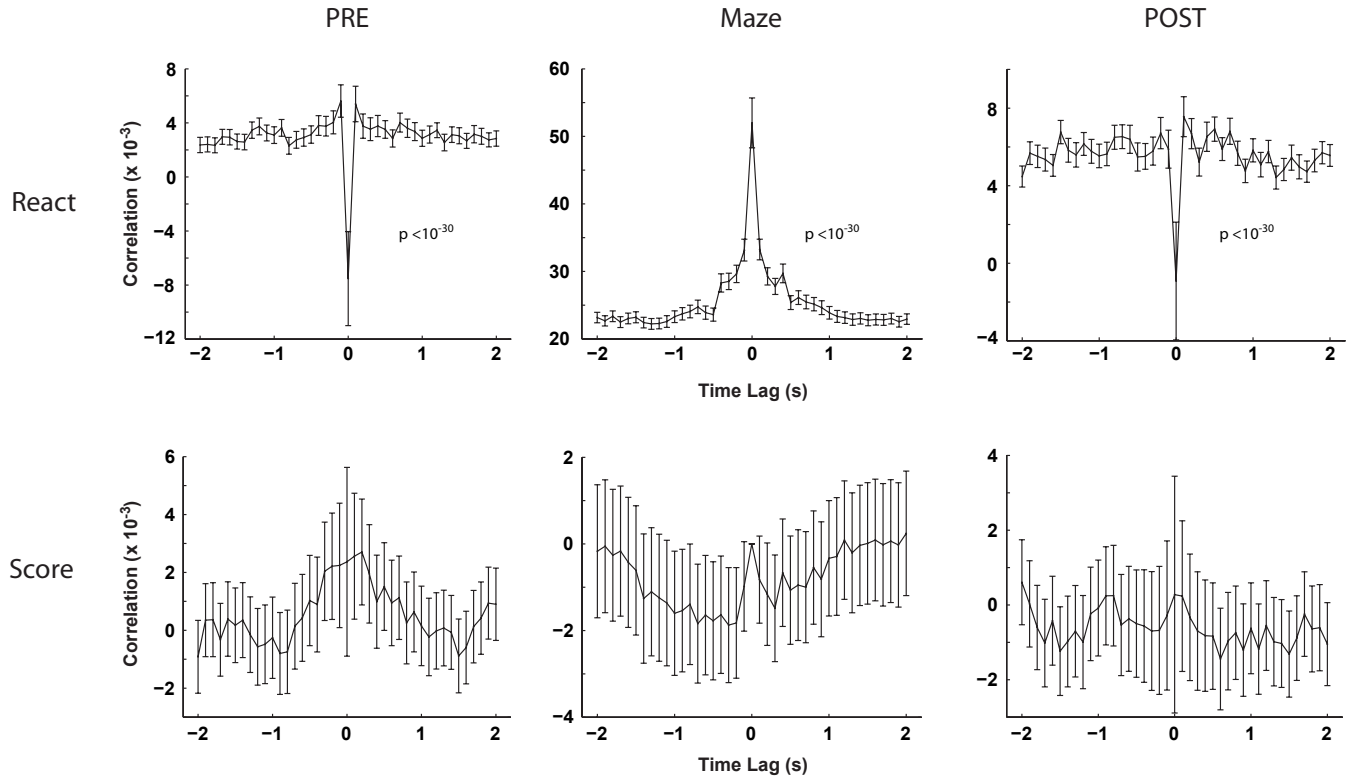

Supplementary Figure 10: Interactions between pairs of signal components. Average cross-correlation between each pairs of significant patterns (averaged over all available pairs) of the reactivation strength (*top*) and the score of the principal components (*bottom*). During sleep epochs, pairs of PCs are more unlikely than chance to co-activate, as in the example in Supplementary Figure 8. This is not a general feature of our reactivation measure, as can be appreciated from the fact that co-activations of pairs of PCs are more likely than chance during the AWAKE epoch. This may reflect some pattern separation phenomenon As expected from the orthogonality of the Principal Components, the cross-correlation for the PC scores is exactly 0 for zero lag during the maze epoch.

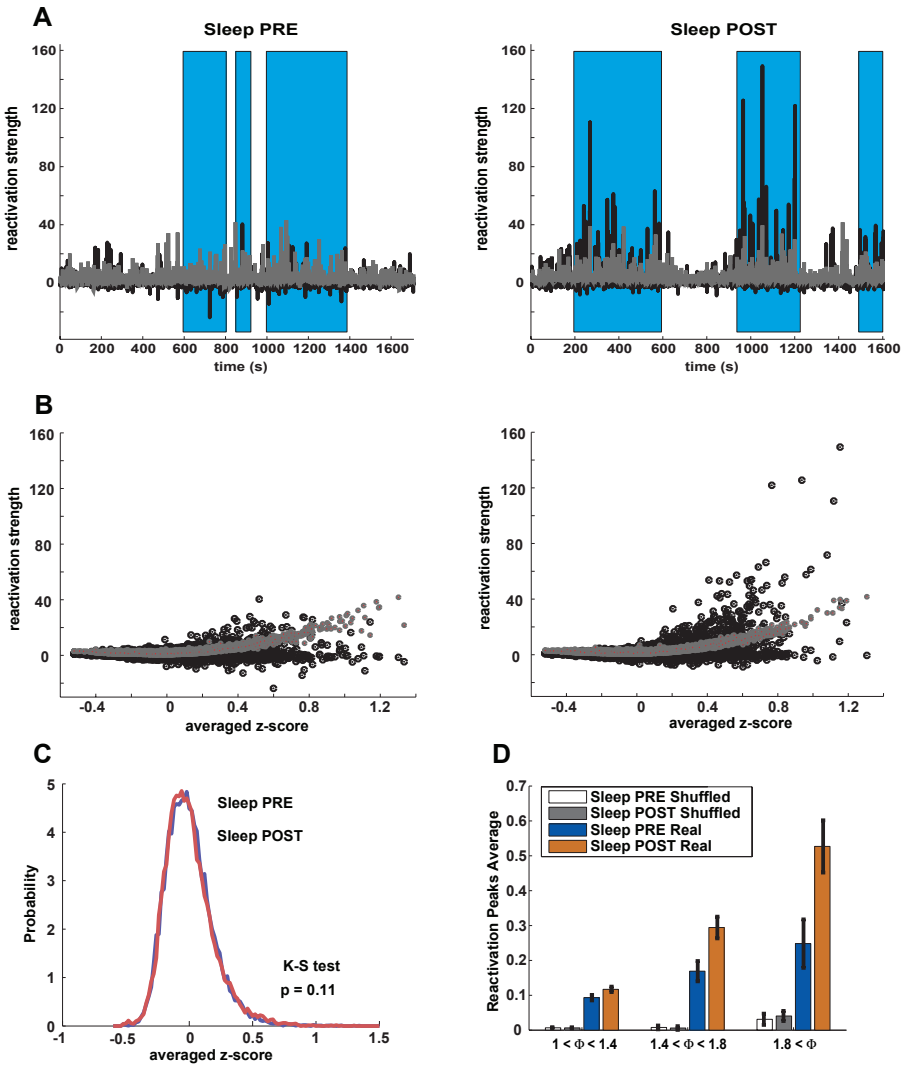

Supplementary Figure 11: Effect of instantaneous global fluctuations of firing rate on the reactivation strength measure. In principle, a greater reactivation strength during POST, could result from an increase in variability of the global firing rate (the average firing rate remaining constant, see Supplementary Fig. 6). In order to control for this possibility, we computed the reactivation strength from shuffled data, that is, at each time bin, the identity of the cells was randomly permuted. This shuffling procedure preserves the instantaneous global firing rate (and its fluctuations), but it destroys the patterns of co-activation. A. for each time bin in an example recording session, the vector of cell activations was shuffled 1000 times, and the reactivation strength was computed from the shuffled vector. The time course of the 95th percentile from the shuffled reactivation strength is shown in gray, while the reactivation strength computed from the actual data is in black. Blue areas denote SWS epochs. B: Scatterplot showing the dependence of the reactivation strength and the instantaneous activation (expressed as the instantaneous z-score averaged over all recorded cells). The actual data are shown in black while the 95th percentile of the shuffled control are in gray, showing that reactivation effects are not likely to be the product of activity fluctuations alone. C: Distribution of the averaged z-score for each time bin in PRE (blue) and POST (red) SWS. The two histograms were not different. These results, taken together, argue against an interpretation of the replay effect solely in terms of a by-product of fluctuations of the global firing rates, suggesting rather that the specific pattern of co-activation is important. D: Average reactivation strength across all analyzed sessions, with shuffled patterns in comparison with actual patterns (as presented in Fig. 2B). For shuffled patterns, the reactivation strength was not significantly different between sleep PRE and POST. Furthermore, it was more than one order of magnitude smaller than for actual patterns.

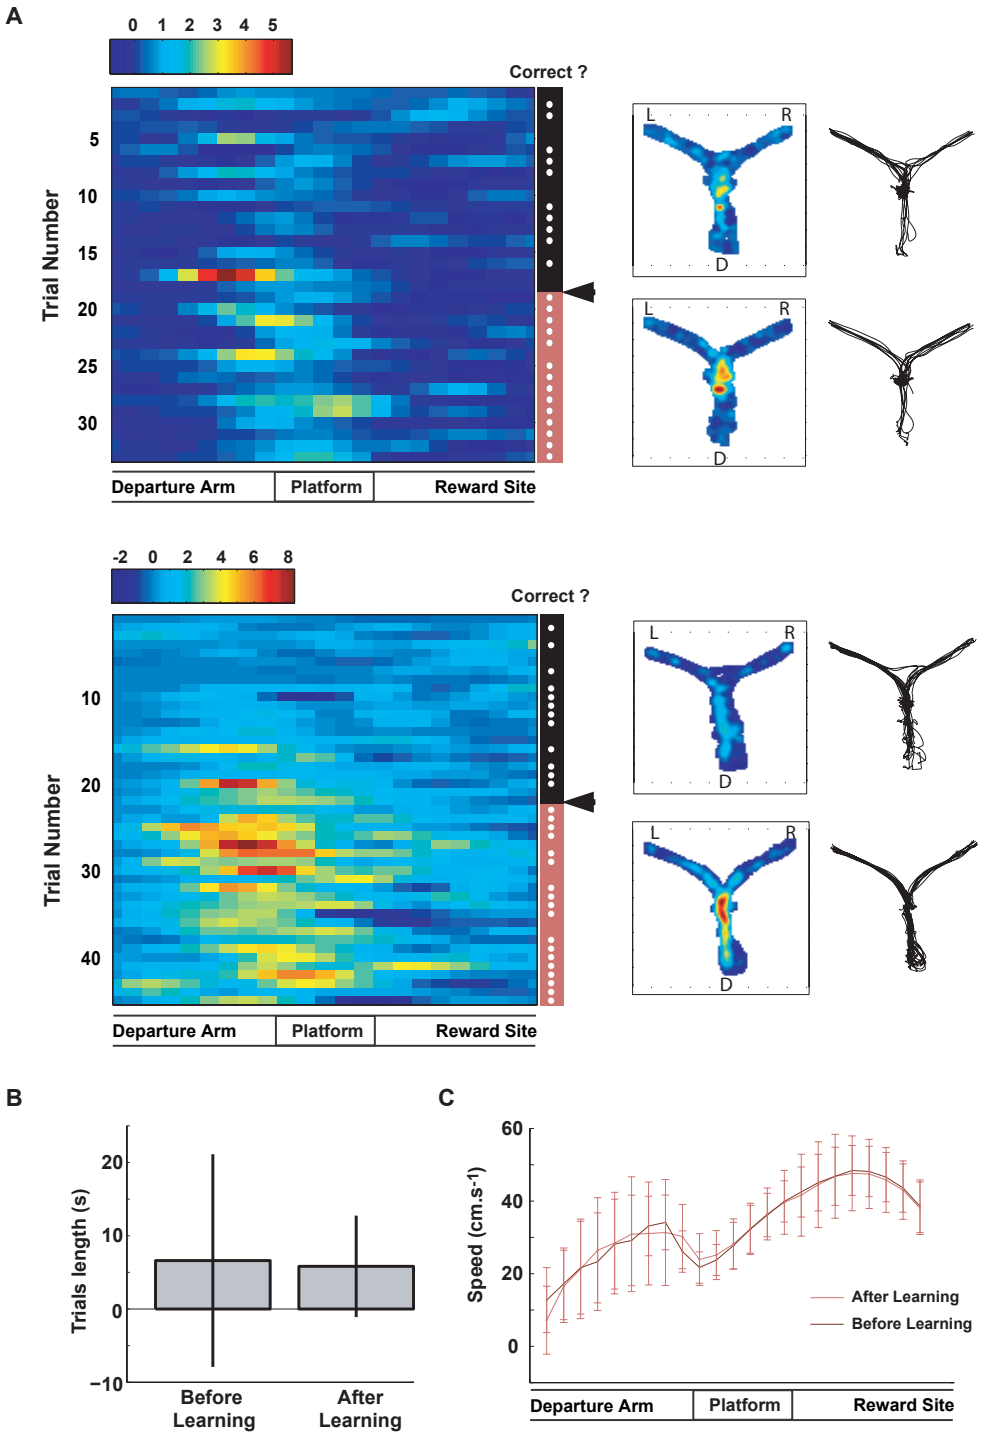

Supplementary Figure 12: Two more examples of spatial distribution of learning-related SPWRs patterns in same format as in Figure 7. Colorbar: reactivation strength. For illustration, the non linearized spatial distribution of reactivation strength (scale not shown) are plotted on the right before (top) and after (bottom) learning trials (D: departure arm; R: right arrival arm; L: left arrival arm). Plots to right show points visited by the rat. B: Average trial length (time from initial movements by rat until arrival at end of arms) before and after learning (10 learning days). The distributions are not significantly different (t-test, $p>0.05$). C: Spatial distribution of animal's speed before (dark red) and after (light red) learning. The speed profiles was not different before and after learning. Thus, at least in terms of gross spatial measures, the behavior of the rat was not dissimilar before and after acquiring the rewarded rule.

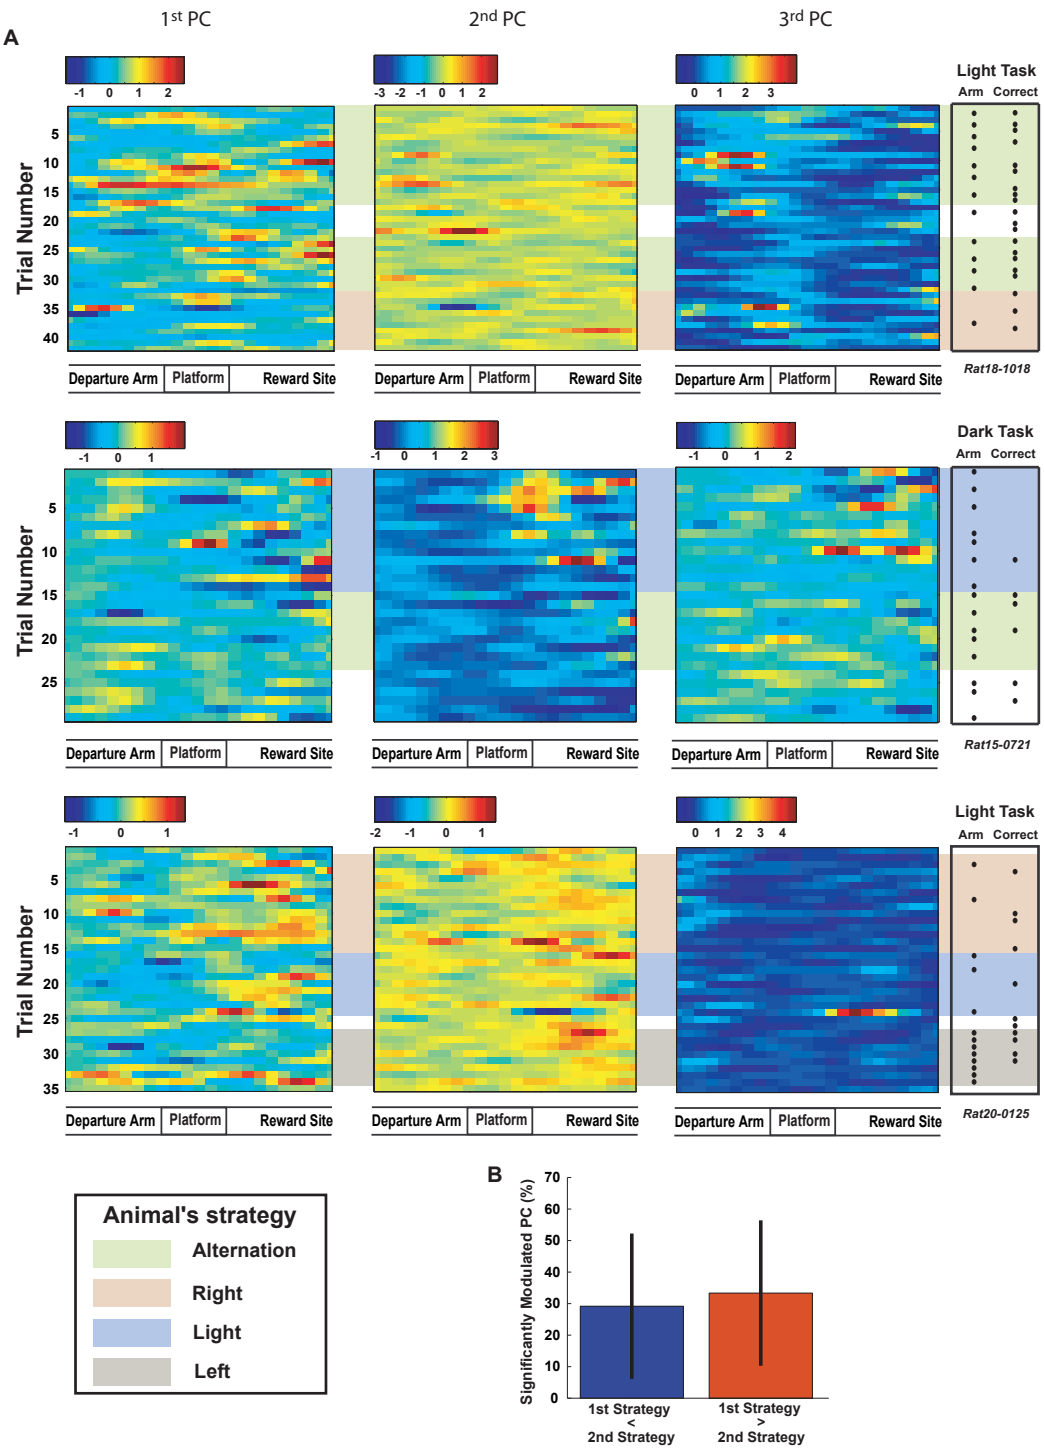

Supplementary Figure 13: A: Spatial distribution of SPWR pattern activations for sessions with shifts between two or more non-optimal strategies. Each sessions are presented in rows, the three first patterns extracted from SPWRs bins are presented for each session in columns. Colorbar: reactivation strength. The panel at the right summarizes the rat behavior: dots on the first column ("arm") denote trial in which the rat selected the left arm, dots on the second column ("correct") denote rewarded trials, background panel color indicates the detected strategy blocks. B: Proportion of patterns significantly stronger for the first strategy or for the second. Differently from what was observed for strategy shifts leading to the rewarded rule, SPWR patterns were equally higher in proportion for the first or the latter strategies.

	# PCs	# Before Learning	# After Learning
SPWR PRE	30	9 (30%)	7 (23%)
SPWR POST	30	1 (3%)	15 (50%)
inter SPWR PRE	42	12 (29%)	10 (24%)
inter SPWR POST	46	15 (33%)	6 (13%)

Supplementary Table 1: Summary Table for Learning Days