

Supporting Information

Functional Coupling Networks Inferred from Prefrontal Cortex Activity Show Experience-Related Effective Plasticity

G. Tavoni, U. Ferrari, F.P. Battaglia, S. Cocco, R. Monasson

I. QUALITY CHECK ON SLEEP EPOCH DURATIONS

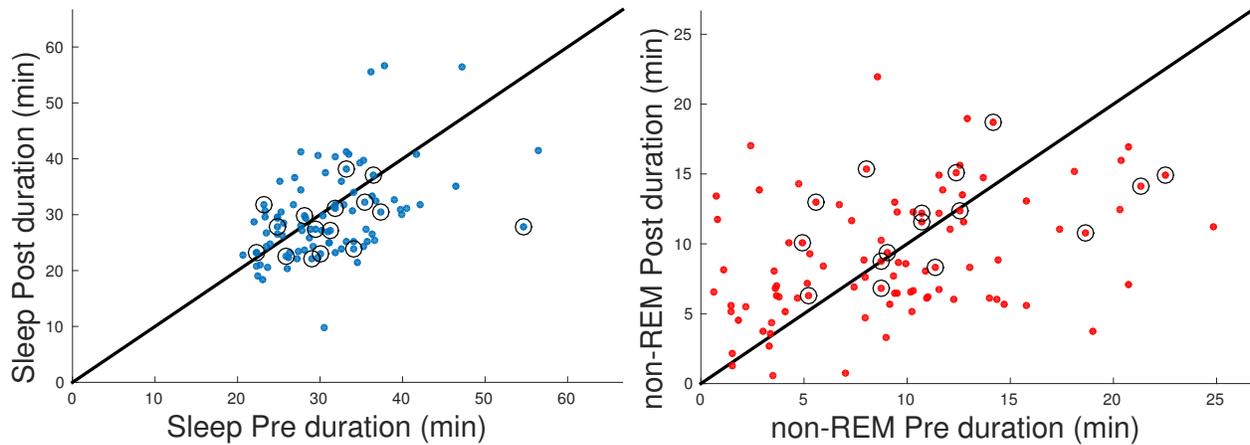


FIG. S1: Quality check on sleep epochs. Left: scatterplot of the epoch duration of Sleep Pre against Sleep Post. Right: scatterplot of non-REM sleep duration in Sleep Pre against Sleep Post. Session showing a potentiated group (see Fig. 5b of the main text) are marked with a circle. The absence of any systematic bias ensures the validity of a comparison among the two sleep epochs.

To ensure the validity of a comparison between the neurons activity in Sleep Pre Task and Sleep Post Task, the two epochs need to have the same quality of sleep. In Fig. S1, respectively, left and right panels, we compare the, respectively, total and non-REM sleep duration for all 96 experimental sessions. As points scatter symmetrically around the diagonal, we exclude the presence of systematic experimental biases, neither in the whole dataset nor in the sessions showing a potentiated group (see Fig. 5b of the main text).

II. DEPENDENCE OF ISING PARAMETERS ON TIME-BIN WIDTH Δt

The inferred local inputs h_i and couplings J_{ij} vary with the value of the time-bin width, Δt . Consider first the case of a single neuron σ , with firing rate f . In the Ising model this neuron is subject to a local input h . The value of the local input is related to the probability $p(\Delta t)$ that the neuron is active in a time bin of width Δt , computed from the spike recordings, through

$$p(\Delta t) \sim f \cdot \Delta t = \frac{e^h}{1 + e^h}. \quad (\text{S1})$$

Hence, we obtain $h \simeq \log(f \Delta t)$ for small Δt : the local input h varies logarithmically with Δt . In a network of more than one neuron with nonzero couplings this calculation does not hold anymore, but Eq. [S1] offers a simple approximation for the dependence of the local input upon the time-bin width.

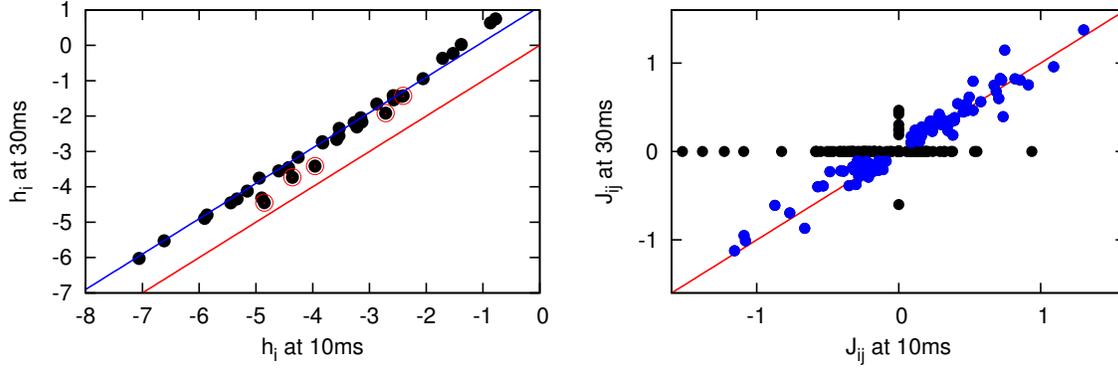


FIG. S2: Scatter plots of the inferred local inputs (left) and couplings (right) for time-bin widths $\Delta t^{(1)} = 10$ ms and $\Delta t^{(2)} = 30$ ms, for the Task epoch of session A. Left: The five red circle correspond to the five neurons in the potentiated group (see main paper). The blue and red lines correspond to, respectively, $h^{30} = h^{10} + \log(\Delta t^{(2)}/\Delta t^{(1)})$ and $h^{30} = h^{10}$. Right: Blue points identify reliable couplings, *i.e.* whose values are larger (in absolute value) than three times their statistical standard deviations. Brown points are unreliable couplings, statistically compatible with zero. The red line corresponds to $J^{30} = J^{10}$. Note that there are more zero couplings for $\Delta t = 30$ ms since the number of neural configurations (time bins) is 3 times smaller than for $\Delta t = 10$ ms: correlation values are known with less accuracy, and can be reproduced within their error bars with a sparser interaction network.

We compare the Ising parameters inferred for two different time-bin widths, $\Delta t^{(1)}$ and $\Delta t^{(2)}$, in Fig. S2. We observe that most local inputs $\{h_i\}$ depend logarithmically on the time-bin width: $h_i^{(2)} - h_i^{(1)} \simeq \log(\Delta t^{(2)}/\Delta t^{(1)})$, see discussion above. This simple rule breaks down for strongly coactivating (coupled) neurons, see Fig. S2, left panel. In addition, the couplings $\{J_{ij}\}$ show no systematic dependence on the time-bin width. Couplings statistically different from zero are, up to small differences, independent of Δt , see Fig. S2, right panel.

III. FOUR-FOLD CROSS-VALIDATION OF THE INFERENCE PROCEDURE

To assess the efficiency of the inference procedure, in the main text Fig 2a&b, we show how the model inferred from 3/4 of the dataset (training set) is able to reproduce the the statistics of the remaining 1/4 (testing set). In Fig. S3, we show the same results for the three other possibilities of dividing the dataset in fourth parts.

IV. COMPARISON OF COUPLINGS AND FIELDS ACROSS EPOCHS

Figure S4 shows the distribution of local input parameters for the three recording epochs, averaged over all 97 sessions. Figure S5 shows the three scatter plots of Ising parameters (local inputs, top row; couplings, bottom row) for each pair of epochs in session A only.

Figure S6 compares the fractions of pairs of neurons in 6 (out of 27) relevant coupling classes with (red) and without (black) conditioning on the Task-epoch class. Conserved ($[-]$ and $[+++]$), potentiated ($[0++]$) and negatively potentiated ($[0-]$) classes contain more pairs than expected by chance.

V. OVERVIEW OF STUDIED SESSIONS

In Fig. S7 we replot the Fig. 3a of the main text on effective potentiation across all session, this time with the labels of relevant sessions. Among the 20 sessions labeled:

- 16 sessions: ABCDEFGHIJKLMNOPQ display a large overall potentiation, as shown in Fig. S7, a large group potentiation and a large difference in the coactivation ratio of the identified potentiated group between the two sleep phases, as shown in Fig. 5b of the main text. These are the 16 sessions on which the study of the reactivation dynamics following ripple inputs has been focused, as shown in Fig. 6 of the main text.
- Sessions H and X are two sessions with large overall potentiations, but in which we do not identify potentiated groups with a large difference in the logarithm of the maximal CoA between the two Sleep epochs. Therefore,

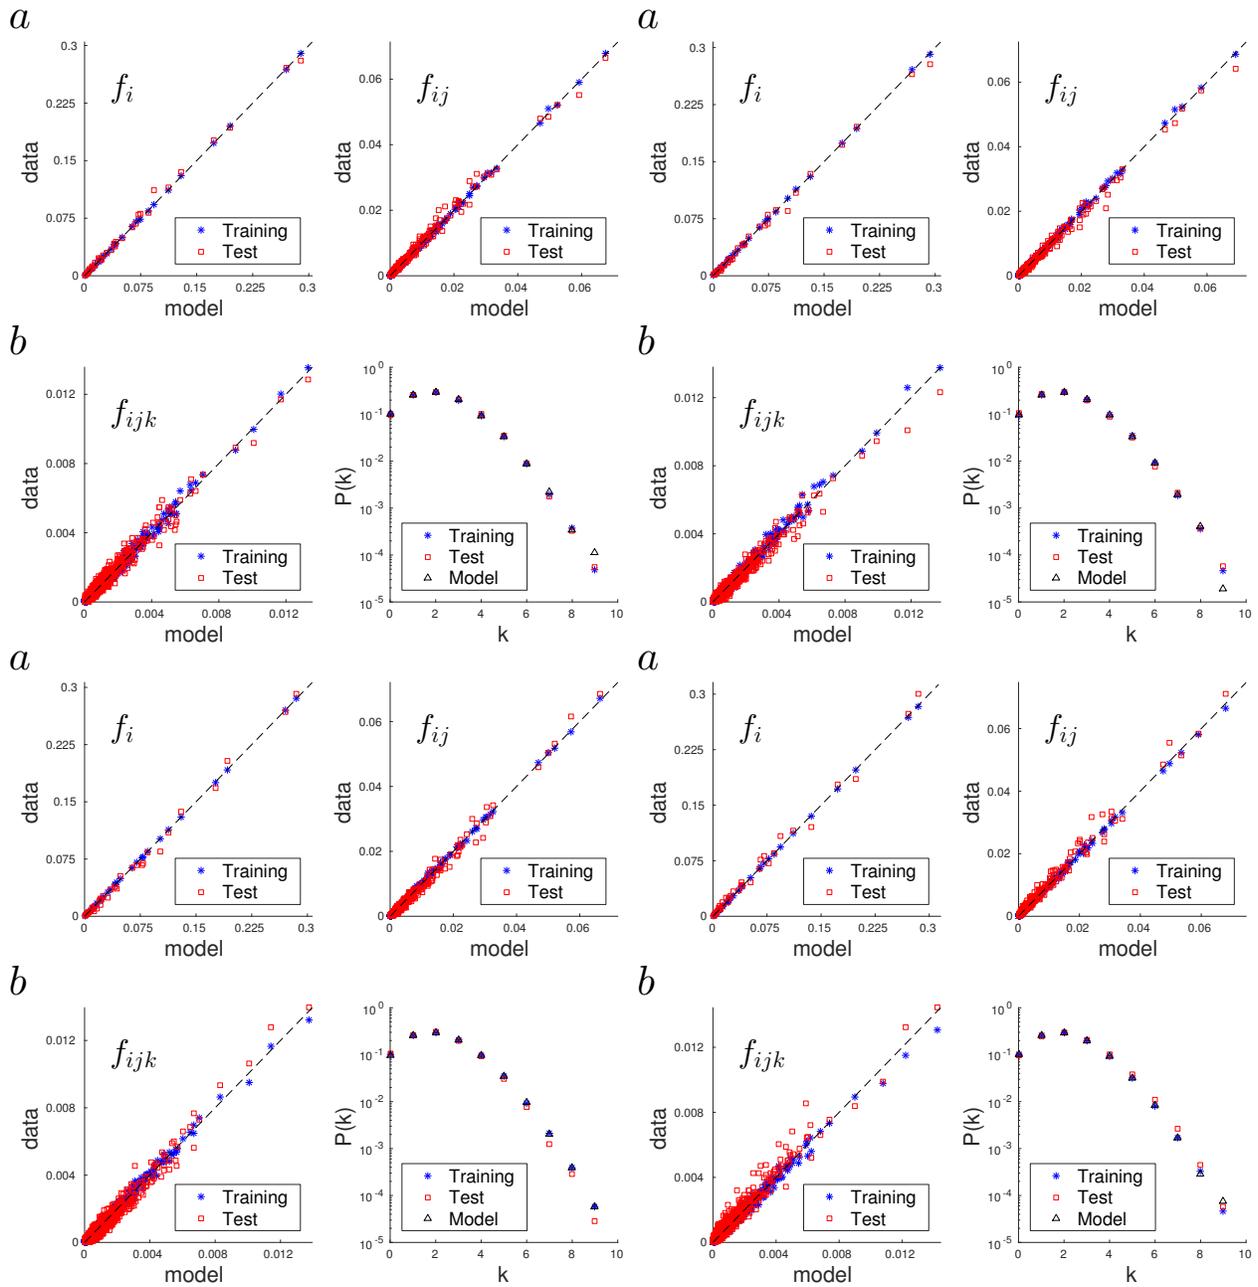


FIG. S3: Four-fold cross-validation of Fig. 2a&b of the main text. To assess the efficiency of our inference procedure, in these plots, we show the same analysis presented in the main text, but applied to the other three subsets generated by the four-fold cross-validation strategy. Model distribution P was inferred from 3/4 of the recorded data and tested on the same data (blue cross) and on the remaining 1/4 of the recording (cross-validation, red squares). For comparison, Figs 2a&b of the main text have been replotted here in the lower right corner.

both sessions were discarded in the analysis of the Ripple-induced reactivation dynamics. Session H has a group potentiation $\simeq 2$ (see Fig. 5b of the main text), and is slightly below the chosen significance threshold having $\Delta \log \text{Max CoA} \approx 1.1$. Session X has also a group potentiation close to 2, but has very low $\Delta \log \text{Max CoA}$ (close to zero, see Fig. 5b of the main text). The disagreement between the group potentiation (about 2, see Fig. 5b of the main text) and the coactivation in session H can be due to sampling issues.

- Finally sessions Y and Z are interesting because they clearly display negative potentiation as we discuss below in more details.

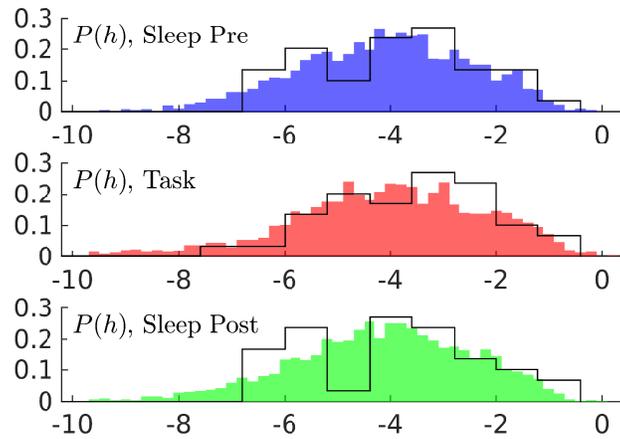


FIG. S4: Distribution of the inferred local inputs h_i in the Sleep Pre, Task and Sleep Post epochs, across all recorded sessions. Black lines show the histograms for session A only.

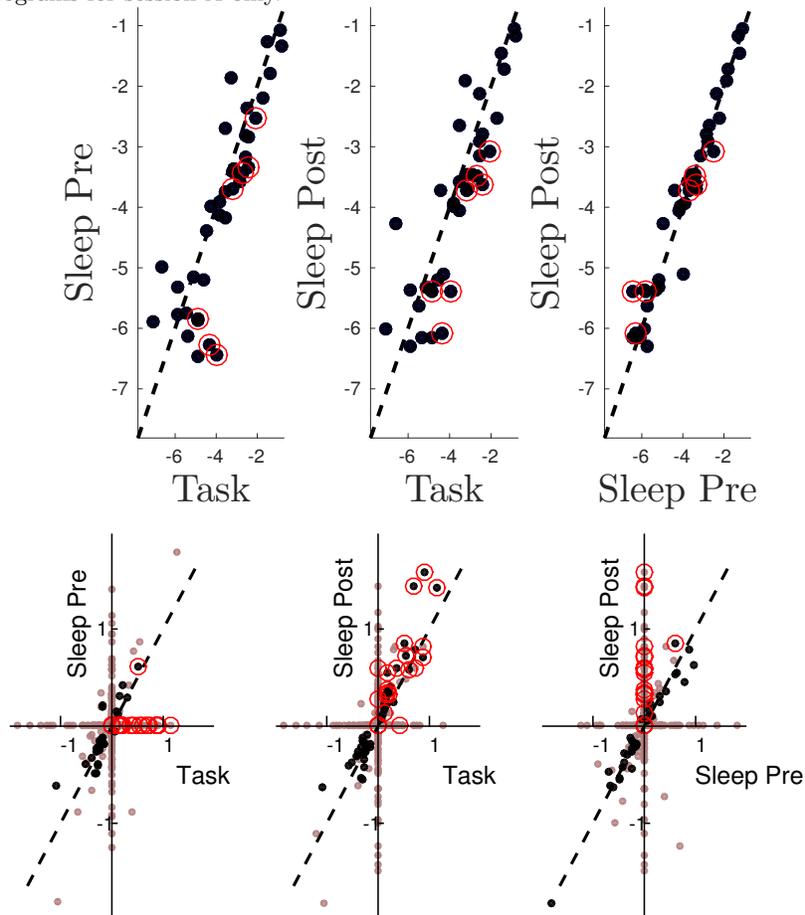


FIG. S5: Scatter of the local inputs (top) and couplings (bottom) across the three epochs of session A.

In the original data sessions are labelled with their letters, and with a 6-digit number RR-MM-DD, where RR stands for the rat number, MM and DD correspond to the month and day of the experiment. Tables 1 and 2 contain the original name of the session, their letter-label, the number of reordered and analyzed neurons and the identified potentiated groups (Table 1) or negatively potentiated groups (Table 2). Note that the data only contain the activity of neurons which, after spike sorting, match between the different epochs. Neurons with very low spiking frequencies have been discarded in the analysis. The list of discarded neurons is in Table. IV, retained neurons are ordered

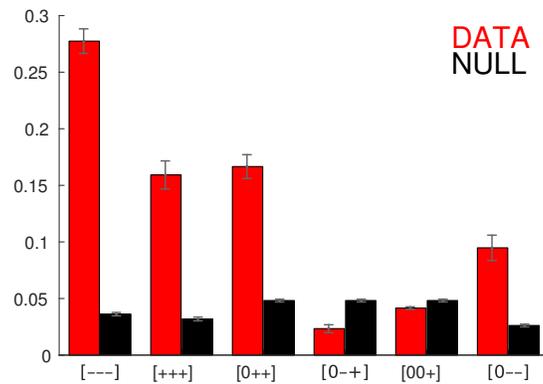


FIG. S6: Bar graph comparing the fraction of pairs of neurons in 6 (out of 27) relevant coupling classes. Red: Task-epoch conditioned fraction, black: null model. Same data as in Fig. 3b of the main text.

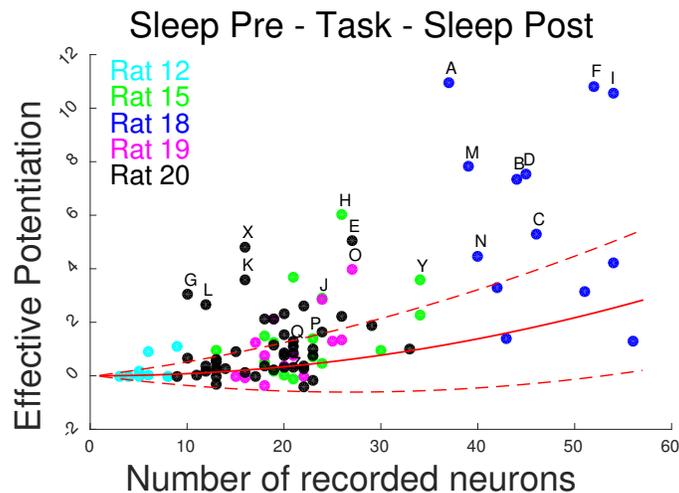


FIG. S7: Effective potentiation on the 97 recorded session, as shown in Fig. 3a of main text, with labels to particular sessions

according to their number.

VI. POTENTIATED GROUPS: IDENTIFICATION, VARIANTS AND STATISTICAL ROBUSTNESS AND SIGNIFICANCE

We have assessed the robustness of the inferred couplings upon sampling by dividing the recording into two halves, and by inferring the model from the two halves separately. For session A the subnetworks supported by the neurons in the potentiated group are very similar, see Fig. S8. In particular, couplings are large and positive in the Task and Sleep Post Task epochs, and strongly increase with respect to Sleep Pre.

Figure S9 shows the top components of the potentiation matrices (see Methods in the main text for definition) of five sessions. Few neurons have large entries exceeding the threshold, and are included in the potentiated groups. In order to identify neurons contributing to the potentiated group, we threshold the entries of the eigenvector of the *Pot* matrix corresponding to the largest eigenvalue. To quantify the impact of different threshold values we compute the p-value for the correlation between the group potentiation and the $\Delta \log \text{Max CoA}$ of the groups (see Fig. 5b of the main paper). Figure S10 shows, as a function of the threshold value c , the behavior of the p-value for all the identified groups (left) and for all groups with at least three neurons (right). We have chosen $c = 0.22$, a value lying in between the two minima (red points).

As shown in Fig. S1, there is an experimental variability in the duration of the sleep epochs, both for the REM and non-REM periods. In order to test the robustness of the effective potentiation against this variability, we have subsampled the longer sleep phases to match both REM and non-REM durations in Pre and Post in new data-sets. We have then reapplied the inference procedure on such subsampled data sets, and recomputed the effective potentiation

Session	Label	# Recorded neurons	Potentiated group
181014	A	37	1 9 20 21 26
181012	B	44	16 21 24 26 31
181026	C	46	2 3 4 5
181021	D	45	9 29 45
201226	E	27	12 15 18 26
181018	F	52	14 26 36 40 43
150726	G	26	3 7 10 26
200208	H	10	4 6 10
181020	I	54	9 22 33
190227	J	24	18 19 23
200118	K	16	1 12 13
200209	L	12	3 7 12
181025	M	39	8 9 35
181011	N	40	3 11 22 29
190228	O	27	7 11
150715	P	23	10 21
200131	Q	21	10 15
200104	X	16	7 11 12

TABLE I: Session, Label, Number of recorded neurons, Potentiated groups found through our spectral method (components of the first eigenvector of the potentiation matrix larger than the fixed threshold 0.22).

Session	Label	# Recorded neurons	negatively potentiated group
150721	Z	34	21 22 26 27
150711	Y	21	10 11 12

TABLE II: Session, Label, Number of recorded neurons, negatively potentiated groups.

and the group potentiation for all sessions. In the left panel of Fig. S11, we compare the resulting effective potentiation averaged over 33 realization of such subsampling with matched durations of the sleep phases with the one obtained from all the data (same as Fig. 3d in the main text). Circles indicate the selected sessions A-Q. No evident bias is seen.

To further test if this variability could affect the potentiation of the identified groups, we have recomputed this group potentiation, for the groups defined in the main text from all the data, over the 33 realizations of the random and uniform sub-sampling of the REM and non-REM sleep periods in Pre and Post. In the right panel of Fig. S11, we compare the sub-sample group potentiation (means and standard deviations are over the 33 realizations) to the one obtained from all the data (same as Fig. 5 of the main text) across all sessions. Even if, for few sessions, we observe a decrease of the group potentiation, results are very robust for a large majority of sessions.

We have tested the CoA of all 160 possible groups obtained from the potentiated group 1-9-20-21-26 of session A upon substitution of one of the five neurons by another neuron among the 37 recorded cells. One group only is found to have a CoA as large as 1-9-20-21-26, both in Task and in non-REM Sleep Post (Fig. S12, cyan curve) and coincides with the variant 1-9-20-26-29. Other two groups among the 160 are found to have large CoA values, comparable to the one of 1-9-20-21-26: these variants are obtained by replacing $26 \rightarrow 10$ (stronger CoA in Task, see Fig. S12, blue curve), and $21 \rightarrow 35$ (stronger CoA in non REM-Sleep Post, see Fig. S12, pink curve). These three variants of the potentiated group are not mutually exclusive: in particular, the CoA of 1-9-20-21-26-29-35 is very large in Task and non REM-Post (Fig. S12, bottom). The CoA of the potentiated groups in the sessions with largest effective potentiations are shown in Fig. S13, together with some of their variants in Fig. S14.

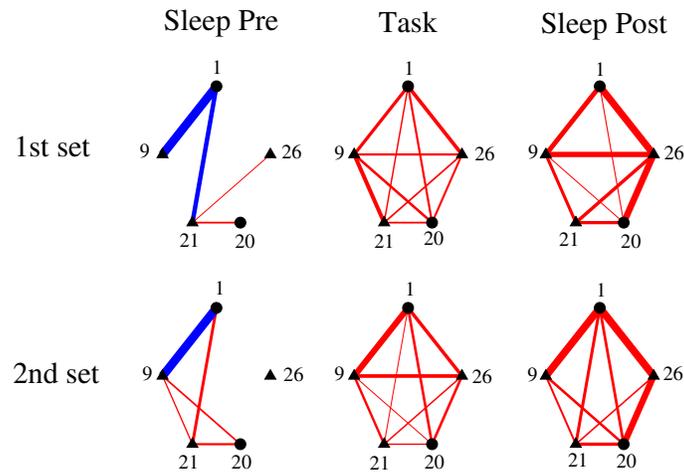


FIG. S8: Subnetworks of couplings over the neurons of the 5-cell potentiated group across the three epochs of session A inferred from the first (top) and the second (bottom) half of the spike-time recording. Blue and red lines refer to, respectively, negative and positive couplings; line thickness is proportional to the absolute value of J . Triangles identify pyramidal cells. This 5-cell group was found with our automatized clustering procedure, and corresponds to the 5 neurons supporting the most potentiated couplings in the 7-cell group shown in Fig. 4b, main text.

Further tests of the statistical significance of the CoAs shown in Fig. 5a of the main text are given in Fig. S15. We compare the CoA of the potentiated group 1-9-20-21-26 with the CoA of groups of 5 randomly chosen neurons (among the $N = 37$ cells of session A). For each random group, we compute the CoA at various time scales τ , and the standard deviation within the Poisson hypothesis of the Methods section in the main paper. The outcome is averaged over 1000 random groups, and shown with black curves in Fig. S15. Note that, as the standard deviation in the case of zero CoA is infinite, we discard those samples in the calculation of the standard deviations; we therefore consider only samples with non-zero CoA, and then multiply the outcome by the fraction of samples with non-zero CoA.

We observe that the CoA of the 5-cell potentiated group is much larger than the CoA for random groups of 5 neurons in the Task and Sleep Post epochs, even when statistical uncertainties are taken into account. This provides further evidence for the statistical significance of the strong coactivation of this group. In Sleep Pre the average CoA is about 2, which shows the existence of a weak positive correlation between randomly chosen cells in the data. The CoA for the 5-cell potentiated group remains equal to zero. This value does not mean that the 5 cells are anti-correlated, and is, indeed, statistically compatible with the independent-cell hypothesis: the product of the spiking frequencies of the 5 cells, multiplied by the duration of the recording is much smaller than one, see Methods in the main paper.

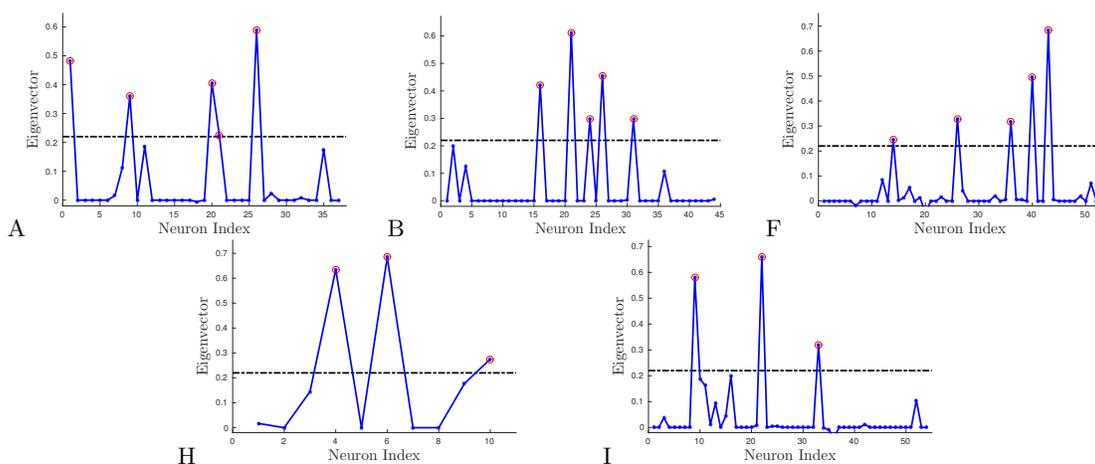


FIG. S9: Top components of the Pot matrix for sessions A, B, F, H, I. The overall sign of \mathbf{v} is chosen so that the majority of entries are positive. Entries larger than the threshold $c = 0.22$ (dashed line) are shown with red circles.

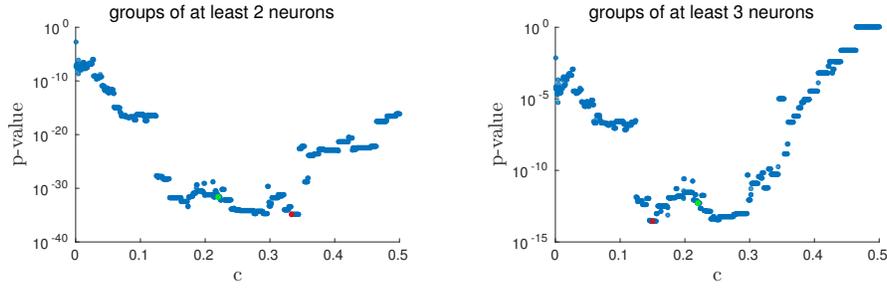


FIG. S10: p -values associated to the linear dependency between $\Delta \log \text{CoA}$ and Pot (see Fig. 5b of the main text), as a function of the threshold c on the entries of the top component for the potentiated group identification. Results are shown for all sessions with potentiated groups with at least 2 (left) or 3 (right) cells. The values of c minimizing the p -value are shown with the red dots; the green dot shows the middle-point $c = 0.22$.

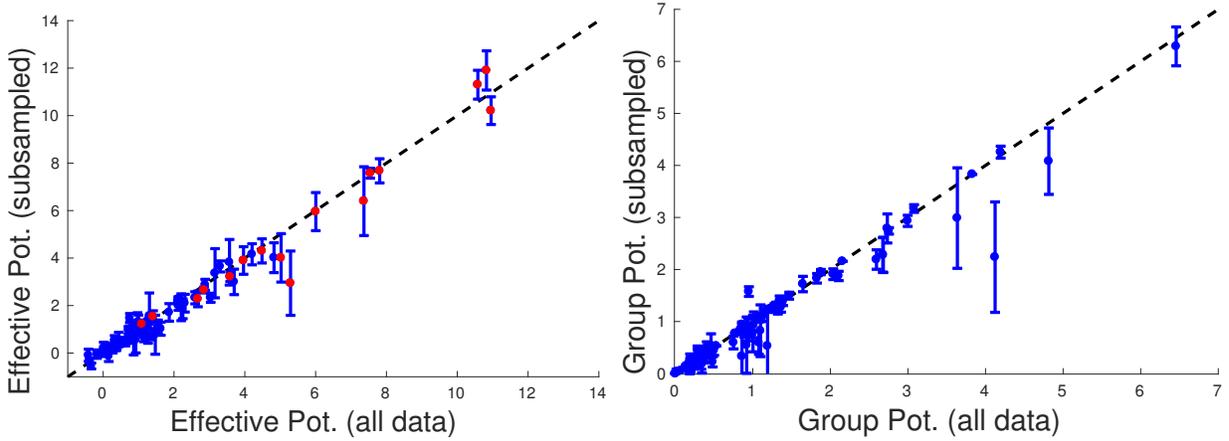


FIG. S11: Left: scatterplot of the session effective potentiation (Eq. (3) and Fig. 3d of the main text), Right: Scatter plot of the mean effective group potentiations (Fig. 5 of main text) from 33 random uniform sub-samplings of the longer sleep epochs, obtained to match the sleep REM and non-REM period durations in Sleep Pre and Sleep Post epochs, against the one obtained from all recorded data. Bars indicate the standard deviations over the 33 sub-samplings. The groups are defined as in the main text, i.e. based on the analysis of the entire recorded data.

VII. NEGATIVE POTENTIATION

In analogy with the effective potentiation in Eqn. [3] of the main text, we introduce the quantity

$$\text{Neg-Pot} = \sum_{\substack{\text{pairs } i, j \text{ with reliable} \\ \text{couplings in Task and Sleep Post}}} \theta(J_{ij}^{\text{Sleep Pre}} - J_{ij}^{\text{Task}}) \times (J_{ij}^{\text{Sleep Pre}} - J_{ij}^{\text{Sleep Post}}), \quad (\text{S2})$$

to measure the effective, session-wide negative potentiation in the functional couplings across the Sleep epochs. Positive contributions to the effective negative potentiation come mainly from class $[0 - -]$ in Fig. 3b in the main text. Figure S16 (left) shows the effective negative potentiation across all sessions. Some sessions, such as Y, display a large effective negative potentiation with respect to a null model obtained through a random reshuffling of the couplings. Other sessions show negative values, mostly coming from the $[+ + +]$ class, with couplings J_{ij} taking values larger in Sleep Post than in Sleep Pre and Task. Similarly to Fig. 3c (right) in the main text we have also swapped the Sleep Post and Sleep Pre data in each session, and re-computed the effective negative potentiation Neg-Pot, with the results shown in Fig. S16 (right). For most sessions, except Y and Z, the available statistics does not allow us to assess that Neg-Pot is significantly different from zero.

VIII. RULES AND BEHAVIORS IN SELECTED SESSIONS

The criterion for rule change and learning point are defined in the main text. In the following we list all sessions studied, describing the learning behavior, sessions are listed in experimental chronological order. When the rule is not changed the changing rule criterium is not reached, unless explicitly stated.

1. 150715:P. Number of trials: 51. The rule is the same as the previous session. Learning point for the first rule: trial 7. Rule change at trial 19. Learning point for the second rule not reached.
2. 150726:G. Number of trials: 34. The rule is the same as the previous session. Learning point not reached.
3. 181011:N. Number of trials: 9. First session of rat 18. Learning point at trial 6. No rule change.
4. 181012:B. Number of trials: 13. The rule is the same as the previous session. Learning point at trial 5. No rule change.
5. 181014:A. Number of trials: 24. The session starts with a new rule with respect to the previous session. Learning point not reached.
6. 181018:F. Number of trials: 42. The rule is the same as the previous session. Learning point not reached.
7. 181020:I. Number of trials: 29. The rule is the same as the previous session. Learning point: trial 10. The changing rule criterion is reached at the end of the session (13 consecutive correct trials) but the rule is not changed in the next session.
8. 181021:D. Number of trials: 30. The rule is the same as the previous session. Learning point: trial 3. It is not surprising that D behaves as a session with consolidate learning, because the changing rule criterion was reached in the previous session even if the rule was not changed by the operator.
9. 181025:M. Number of trials: 20. The rule is the same as the previous session. Learning point: trial 1.
10. 181026:C. Number of trials: 34. The rule is the same as the previous session. Learning point for the first rule: trial 1. Rule change at trial 11. Learning point for the second rule not reached.
11. 190227:J. Number of trials: 20. The rule is the same as the previous session. Learning point: trial 16.

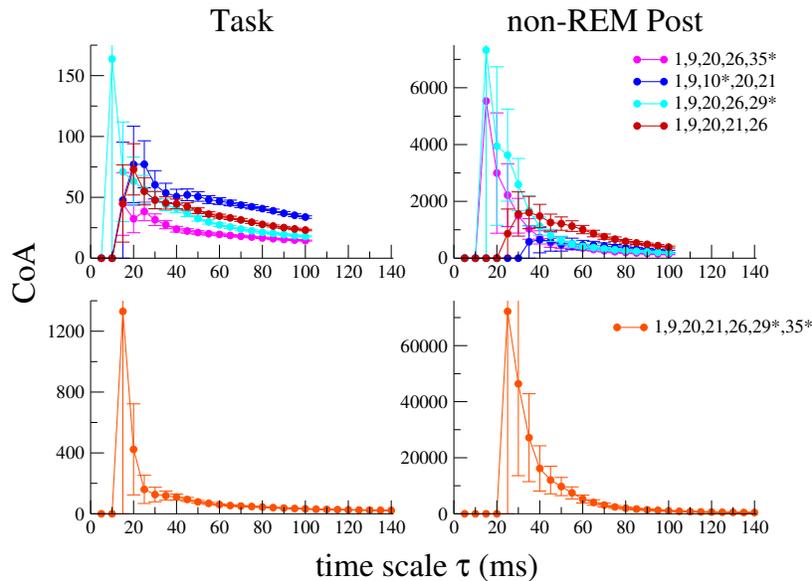


FIG. S12: Top panels: CoA of the potentiated group 1-9-20-21-26 and of the three groups with largest CoA obtained upon substitution of one neuron (indicated by stars). Note the large error bars on the CoA of the cyan and magenta group at small time scales. Bottom panels: CoA of the 7-neuron group, which extends the potentiated group 1-9-20-21-26 by adding neurons 29 and 35.

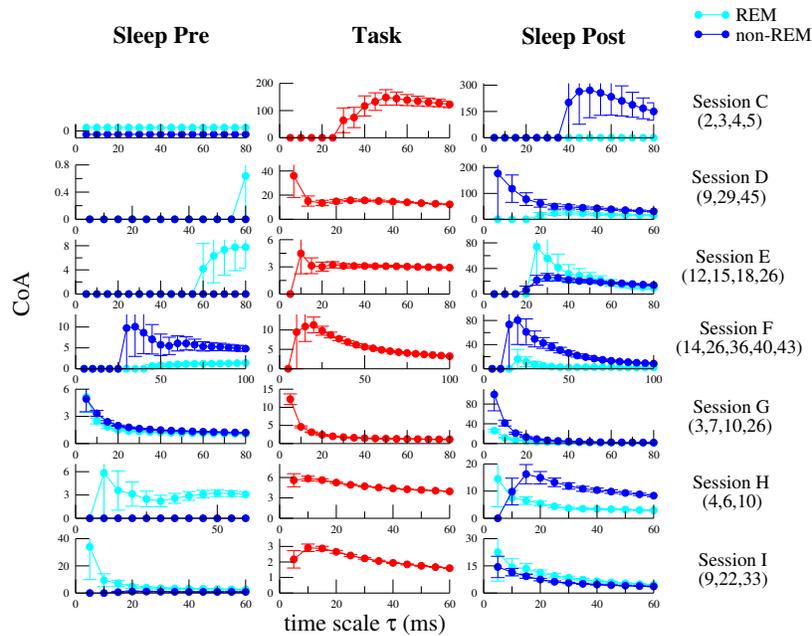


FIG. S13: CoA in the different epochs of potentiated groups for the sessions with largest effective potentiations. The potentiated groups are obtained with fixed threshold $c = 0.22$; Note that in sessions D-H-I where the rat has learned the rule a large CoA is present also in the REM phase, while in session F and G, where the rat has not learned the rule, CoA is substantially larger in non-REM than in REM, and in session C, where the rule has been shifted, CoA is only present in the non-REM phase. In all the sessions above the rule was the same of the previous session and in session E, F, G the CoA of the potentiated group is also present in the Sleep-Pre epoch.

12. 190228:O. Number of trials: 15. The rule is the same as the previous session. Learning point: trial 8. Note that J and O were recorded on the third and fourth days of the same rule, in the same rat. On day 1 the rat had almost reached the learning point (6 correct trials and two errors at the end of the session), then showed poor performances for the next day, it reached the learning point at the end of the session J on day 3 (5 correct trials), and consolidated the learning on day 4 (O) (8 correct trials at the end) but did not reach the rule-changing criterion. It is therefore not surprising that O behaves as sessions with consolidated learning in Fig. 6c (main text).
13. 200118:K. Number of trials: 34. Rule change at trial 18 (but no rule change from the previous session to the beginning of this session). Learning point for the first rule: trial 6. Second rule is not learned.
14. 200131:Q. Number of trials: 36. Rule change at trial 25 (but no rule change from the previous session to the

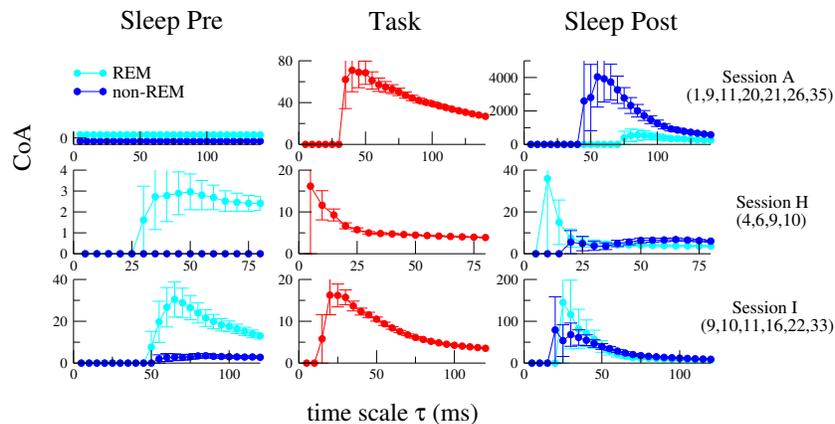


FIG. S14: CoA in the different epochs of potentiated groups for the sessions with largest potentiation obtained with a slightly smaller value of the threshold than $c = 0.22$.

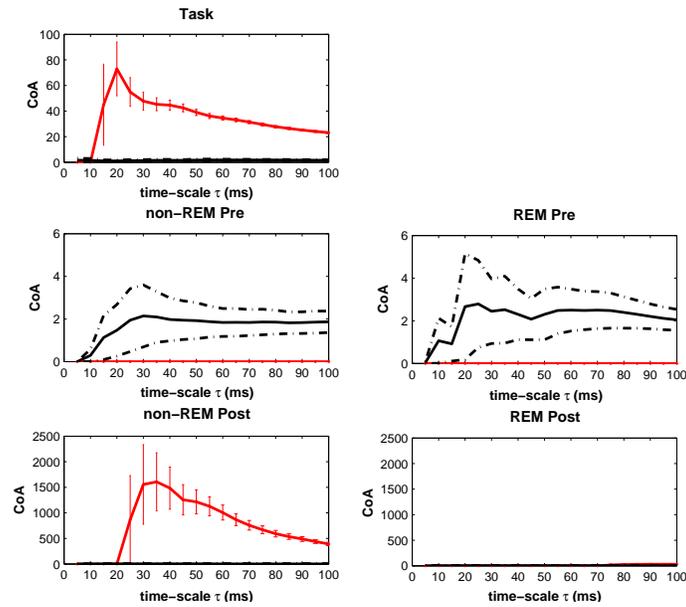


FIG. S15: Coactivation factors for groups of five randomly chosen neurons in Session A (black curves: average of 1000 samples (full), \pm average over the standard deviations (dashed-dotted)), compared to the CoA for the 5-cell potentiated group (red curves, corresponding to the results shown in Fig. 5a of the main text).

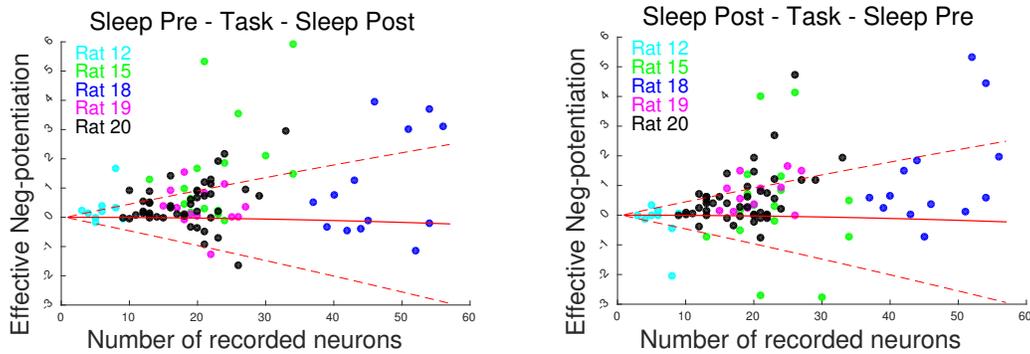


FIG. S16: Left: Negative potentiation, see Eq. [S2], across all the sessions. Right: Control obtained by swapping the Sleep Pre and Sleep Post epochs.

beginning of this session). Learning point for the first rule: trial 5. Second rule is almost learned: 6 correct, 1 error, 1 correct at the end, which means performance at exactly 80% after 3 consecutive correct trials.

15. 200208:H. Number of trials: 35. The rule is the same as the previous session. Learning point: trial 30.
16. 200209:L. Number of trials: 34. 2 rule changes at trials 13 and 29 (but no rule change from the previous session to the beginning of this session). Learning point for the first rule: trial 3. Learning point for the second rule: trial 13 (first trial with the second rule). Third rule is not learned.
17. 201226:E. Number of trials: 39. No rule change within the session or with respect to the previous one. No learning.

IX. RIPPLE-RESPONSE AND LEARNING BEHAVIORS IN THE 17 SELECTED SESSIONS

Table III summarizes the results on Ripple-conditioned Reactivation (*RR*) of Fig. 6b (main text) and on Auto-Reactivation (*AR*) in Fig. 6c (main text) of the potentiated group, and correlate them session by session with the learning behavior, which is detailed in the previous section.

Session	Learning	Fast RR	Slow RR	AR Amplitude
F	NO	YES	NO	Low
G	NO	YES	NO	Low
A	NO	YES	NO	Medium
E	NO	YES	NO	Medium
H	YES(1)	YES	NO	Medium
M	YES(1)	YES	NO	Medium
B	YES (1)	YES	YES	Medium
I	YES (1)	YES	YES	Medium
N	YES(1)	YES	YES	Medium
J	YES(1)	NO	NO	Medium
O	YES(1)*	NO	NO	Large
D	YES(2)	NO	NO	Large
L	YES(2)-YES(2)-NO	NO	NO	Medium
C	YES(2)-NO	NO	NO	Large
P	YES(2)-NO	NO	NO	Large
K	YES(2)-NO	Not sampled	Not sampled	Large
Q	YES(2)-YES(1)	YES	YES	Large

TABLE III: Overview of the Ripple-conditioned Reactivation (RR) responses and Auto Reactivation (AR) and learning behavior. Sessions are ordered following the level of learning: no learning: NO, learning point reached : YES(1), learning point & changing rule criterion reached: YES(2). As described in the previous section in session O the change rule criterion is not reached, but long-lasting learning is present. In the change rule case the Table also specifies if the following rule is learned (e.g. sessions L and Q) or not (e.g. session C). In Session L two rules are learned and changed. Fast and Slow components to RR are reported as present (YES) or absent (NO) if their z-score amplitudes are ≥ 2 or ≤ 2 (Fig. 6b Main Text); AR Amplitude (Fig. 6c Main Text) is classified depending on its value at slow time scales as Low (≤ 2), Medium (≥ 2 and ≤ 8) and Large (≥ 8).

[1] Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I. and Battaglia, F.P., Replay of rule-learning related neural patterns in the prefrontal cortex during sleep, *Nature Neurosci.* **12**, 919-26 (2009).

Session	Removed neurons	Session	Removed neurons
150628	14	200126	6
150629	13	200130	1
150705	12	200206	2
150713	23	200208	4
150721	21	200209	2
181017	33	200210	12
181019	4, 21, 55	200214	16, 22, 24
181020	24	200216	11
181102	6	200222	8, 9, 11, 13
181103	24	200223	9, 13
190224	4, 17	200224	7, 9
190301	2, 18	200227	6
190303	2, 17	200228	8
190308	9, 15	200301	11, 13, 15, 20, 26, 28, 33, 36, 37, 38, 42, 49
190313	12	200303	9, 10, 13
190314	11	200308	10, 11, 12, 13
200102	16	200309	8
200103	1	201219	1, 2, 3, 4, 5, 6, 8, 9, 10
200104	2, 16	201220	6, 25, 26
200107	17	201221	4, 5, 25
200109	1, 13	201222	13, 17, 19, 20, 24
200110	17	201223	6, 24
200111	1, 11	201226	19
200112	3	201228	23
200116	3, 6, 8, 22	201229	17, 18
200124	4	201230	11
200125	5		

TABLE IV: Table of neurons removed from the data [1], spiking less than 10 times in each epoch