

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

Gaia Tavoni ^{1,2†}, Ulisse Ferrari ^{1,2†}, Francesco P. Battaglia³, Simona Cocco ^{1*}, Rémi Monasson ²

1 Laboratoire de Physique Statistique de l'Ecole Normale Supérieure - UMR 8550, associé au CNRS et à l'Université Pierre et Marie Curie, 24 rue Lhomond, 75005 Paris, France

2 Laboratoire de Physique Théorique de l'Ecole Normale Supérieure - UMR 8549, associé au CNRS et à l'Université Pierre et Marie Curie, 24 rue Lhomond, 75005 Paris, France

3 Donders Institute for Brain, Cognition and Behaviour, Radboud Universiteit, Nijmegen, The Netherlands.

† these authors contributed equally to the work;

* corresponding author: cocco@lps.ens.fr

Abstract

We infer functional coupling network models reproducing the spiking activity of simultaneously recorded neurons in prefrontal cortex (PFC) of rats, during the performance of a cross-modal rule shift task (learning epoch), and during preceding and following sleep epochs. A large-scale study of the 97 recorded sessions allows us to detect, in ~20% of sessions, effective plasticity between the sleep epochs. The functional couplings may either increase (effective potentiation) or decrease (effective depotentiation) across the sleep epochs, depending on their values in the learning epoch. Moreover these coupling modifications are supported by a small subset of the recorded neurons, which we identify by means of an automatized procedure. We verify that the potentiated groups of neurons undergo an increase in their frequency of coactivation in the spiking data between the two sleep epochs, and, hence, participate to putative learning-related cell assemblies. Study of the reactivation dynamics of the potentiated groups suggests a possible connection with the learning behavior. Reactivation is largely driven by hippocampal ripple events when the rule is not yet learned, and is much more autonomous, and presumably sustained by the potentiated PFC network when learning is consolidated. These results highlight the capability of inference approaches to characterize by means of functional connectivity experience-related effects in population neural recordings.

Author Summary

Cell assemblies, *i.e.* groups of coactivating neurons coding for memories, are widely believed to emerge through modifications of interactions between neurons *e.g.* resulting from learning. We propose an approach to identify learning related cell assemblies from multi-electrode recordings of neural activity *in vivo*, and apply it to the prefrontal cortex (PFC) activity of behaving rats. Our statistical physics inspired approach consists in inferring effective, functional interactions between the recorded cells, which reproduce the correlations in their spiking activities. The analysis of the effective interaction networks allows us to identify changes (effective potentiation or depotentiation) in the functional interactions between the sleep epochs preceding and following the learning of a task, which are correlated to the functional network associated to the neural activity during learning. We find that the couplings undergoing effective

potentiation are supported by a subgroup of the recorded cells. These neurons, which coactivate during learning more than expected by chance, increase their co-activation frequency in the sleep epoch after learning compared to the sleep epoch before learning, and, hence, belong to an experience-related cell assembly. Coactivation of this cell assembly can be explained as a fast response to hippocampal inputs (ripples), known to be important for memory consolidation, in some sessions where the rule has not been learned yet. In many other sessions in which the rule has been learned, however, reactivation of the cell assembly show a strong slow dynamical component, often unrelated to ripples, which presumably reflects the existence of a potentiated PFC synaptic network.

Introduction

In recent years, many techniques have been developed to monitor brain activity in a detailed way and provide a multi-channel, multi-dimensional picture. Different channels may represent coarse-grained activity of neurons in $\sim 1 \text{ mm}^3$ volumes like in functional magnetic resonance imaging (fMRI) [1], or give access to the activity of single neurons, such as with multi-electrode electrophysiological recording [2–4] or calcium imaging [5–7]. The development of these techniques was motivated by the need to characterize the brain state at the network level, and to understand how connections between neurons determine the dynamics and the information processing of neural ensembles. While much progress has been made in characterizing the *connectome*, *i.e.* the exact pattern of connections between neurons in a brain circuit [8], the link between this anatomical architecture and neural network function remains elusive, at least for large-scale circuits.

A useful notion in this context is that of *functional couplings*. Functional couplings are generally derived by a reverse inference procedure [9]: connections between neurons are assumed to determine neural activity, by way of a statistical or dynamical model, and the coupling values are calculated as those most likely to produce the observed activity data. Compared to earlier estimates of functional interactions based on activity correlations [10–13], reverse inference techniques have the advantage of discounting correlations due to interaction paths going through third-party neurons between the recorded cells, therefore providing a much sharper picture of the underlying interactions. Furthermore, functional coupling inference provides a helpful metaphor of the studied neural circuit, under the form of an artificial network with a connectivity matrix determined from the data, and whose properties may be studied with analytical or numerical methods, helping to understand how neural computations, learning and memory work.

Functional coupling can be compared with anatomical couplings, when available, therefore contributing to understand how structural connections shape neural activity. While functional couplings are not likely to match *one-on-one* structural couplings, the fingerprint of neural interactions they provide could be used to track changes in the underlying connectivity. A central paradigm in neuroscience is that couplings are plastic: learning and memory consolidation happen through changes in neural synaptic couplings following repeated coincident pre and post synaptic activations, as was postulated by D. Hebb and proven by *in vitro* experiments showing Long-Term-Potentiation (LTP) and Long-Term Depression (LTD) after repeated coincident and non-coincident stimulations [14, 15].

Hebb further postulated that cell assemblies [16–18], closely connected, synchronously activating groups of cells are the main constituents of memory and information representations. The activation and reactivation ('replay') of cell assemblies is thought to be critical for consolidation and re-elaboration of memories, working memory and decision making [19–21]. The precise characterization of cell assemblies from experimental data remains, however, very difficult. Current available methods for cell assembly detection and replay estimation often rely on the identification and on the matching of templates [22–24]. In the hippocampus, for instance, such templates are provided by the temporal sequence of firing events of place cells during the awake phase. The pairwise cross-correlation matrix can also be used to search for clusters of neurons with related firing patterns [25, 26], or to approximate templates from principal

component analysis (PCA) [27, 28].

The latter approach was used in particular to analyze the prefrontal cortex activity of behaving rats, recorded during the awake epoch and during the preceding and subsequent sleep phases [27]. During the awake epoch the animal faced a learning task, where it had to find the rewarded arm in a Y-shaped maze; this arm was chosen according to a rule (left or right arm, or where the light is on or off) set by the operator. As soon as the rule was consistently learned it was changed. PCA-based analysis of the recorded activity showed that the activity of the learning phase was replayed during the subsequent sleep; this replay is at the basis of memory consolidation.

Here, we re-analyse the same recordings, and find the patterns of functional couplings by inferring graphical models expressing the conditional dependencies between the spiking events of the recorded cells. We make use of the maximum-entropy Ising model from statistical mechanics, which parameters are tuned to reproduce the recorded firing frequencies and pairwise cross-correlations [29]. This inference approach has been tested on several multi-electrode recording of both *in vitro* and *in vivo* neural activity [29–32]. Working with the inferred couplings rather than considering directly the *correlations* in the data, allows us to refine the analysis of the recording and to unveil modifications in the functional couplings between the two sleep epochs (effective potentiation or depotentiation), which are consistent with the functional network derived from the learning epoch. Effective potentiation is supported by a subset of the recorded cells, which we identify by means of an automatized procedure. Our findings are supported by a large-scale study of about 100 experimental sessions. Despite the variations from session to session, presumably due to the partial and random sampling of cells, we are able to identify in about 20% of sessions a potentiated group. We then investigate in the data the change in the collective firing properties of the identified potentiated groups, and find a strong increase in coactivation for such groups, likely to belong to a learning-related cell assembly. Analysis of the reactivation dynamics of the potentiated group suggests a direct relationship with the learning behavior of the rat. In sessions where the rule has not been learned yet the reactivation can be essentially explained as a fast response to hippocampal inputs (ripples), known to be important for memory consolidation. In many other sessions in which the rule has been learned, however, the reactivation shows a strong slow dynamical component, often unrelated to ripples, which presumably reflects the existence of a potentiated PFC synaptic network. Besides the application to a specific set of prefrontal cortex data considered here we think that our approach could be of general interest to identify the presence of subtle changes in the correlation structure of population recordings, and to understand their relation with learning, one central topic in Neurosciences [18].

Results

We have re-analyzed recordings of the activity of tens of neurons in the prefrontal cortex of five behaving rats [27], see Material and Methods for more details. Each one of the 97 recording sessions is divided into three \simeq 30-minute epochs: a Task epoch in which the rat had to learn a cross-modal rule (go left, right, where the light is on, or off, in a Y-shaped maze), which was changed as soon as the rat had learned it, and two Sleep epochs, one before (Sleep Pre) and one after (Sleep Post) the Task epoch. Through spike sorting one can identify the same neurons recorded in the different epochs (Sleep Pre, Task, Sleep Post) of a session; the number N of neurons reliably mapped in all three epochs varies from 3 to 56 depending on the session. No mapping could be established between different sessions.

Inference of functional coupling networks.

We briefly present the approach to model the distribution of activity of the N recorded neurons, see Methods section. The spiking times are binned within small time bins of duration $\Delta t = 10$ ms, as illustrated in Fig. 1a; see Supporting Information-I for a discussion of the time-bin choice. The activity configuration $(\sigma_1, \sigma_2, \dots, \sigma_N)$ is a snapshot of the neural activity, where σ_i takes values one or zero depending on whether the i -th neuron is, respectively, active or inactive in the time bin . We define f_i and f_{ij} as the average values over time bins of, respectively, σ_i and $\sigma_i \sigma_j$: f_i represents the probability that neuron i is active in a time bin, and f_{ij} denotes the joint probability that both cells i and j are active in a time bin.

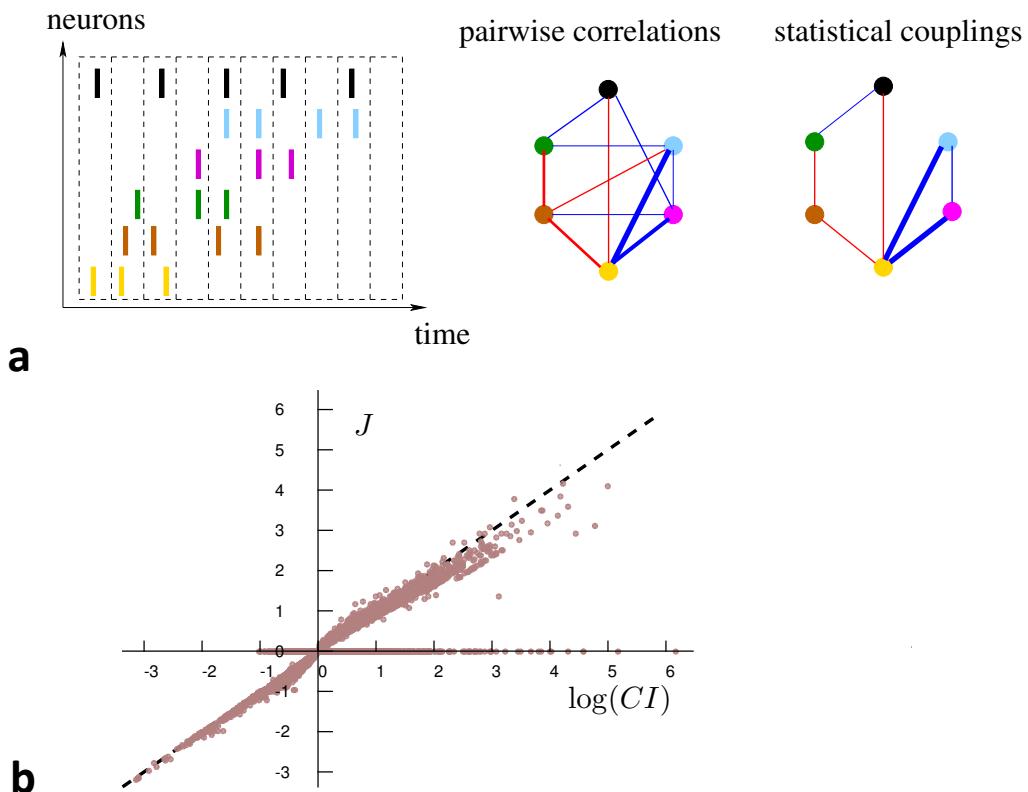


Figure 1. From spiking data to the Ising model. **a.** Spiking times are binned into time-bins of width Δt ; each neuron i is assigned the variable $\sigma_i = 1$ or 0, depending on whether it is active or not in the time-bin (left). Pairwise correlation indices CI_{ij} are computed from this binned data, and define a network of correlations (middle). The network of statistical couplings J_{ij} defining the Ising model distribution P , Eqn. [1], is generally sparser (right). Red and blue links correspond, respectively, to $CI > 1$, $J > 0$ and to $CI < 1$, $J < 0$; the widths are proportional to the absolute values; Links corresponding to CI or J smaller than one tenth of the maximal correlation index or coupling are not shown. **b.** Scatter plot of the inferred couplings J_{ij} vs. logarithm of the correlation indices CI_{ij} for all pairs of cells, and all epochs and sessions. For many pairs (i, j) both quantities are very similar (Methods). However, we observe that many other large correlation indices may be explained by network effects, and correspond to vanishing couplings. The inferred coupling network is therefore much sparser than the correlation network, see Fig. 1a.

We model the probability distribution of activity configurations as

$$P(\sigma_1, \sigma_2, \dots, \sigma_N) = \frac{1}{Z} \exp \left(\sum_{i < j} J_{ij} \sigma_i \sigma_j + \sum_i h_i \sigma_i \right), \quad [1]$$

where Z ensures normalization of the distribution. P in Eqn. [1], called Ising model in statistical physics, is the least constrained (with maximum entropy), default probability distribution reproducing this low-order spiking statistics [29]. We look for the Ising model parameters $\{h_i, J_{ij}\}$ such that f_i and f_{ij} match, respectively, the average values of σ_i (for all neurons i) and $\sigma_i \sigma_j$ (for all pairs of neurons i, j) over P . To do so we use the Adaptive Cluster Expansion (ACE) inference algorithm [32–34], which also gives access to the statistical uncertainties $\{\delta h_i, \delta J_{ij}\}$ over the inferred parameters (Methods). Parameters h_i define effective local inputs which tune neuronal frequencies. Parameters J_{ij} define the effective pairwise couplings between the cells (Fig. 1a): J_{ij} different from zero expresses the presence of a conditional dependence between neurons i and j , not mediated by other neurons in the recorded population. The conditional average activity of neuron i given the other neuron activities $\{\sigma_j\}$, with $j \neq i$, reads

$$\langle \sigma_i \rangle = \frac{P(\sigma_1, \dots, \sigma_i=1, \dots, \sigma_N)}{P(\sigma_1, \dots, \sigma_i=0, \dots, \sigma_N) + P(\sigma_1, \dots, \sigma_i=1, \dots, \sigma_N)} = \frac{e^{V_i}}{1 + e^{V_i}}, \quad \text{with } V_i \equiv \sum_{j(\neq i)} J_{ij} \sigma_j + h_i. \quad [2]$$

It is a logistic function of its total input, V_i , equal to the sum of the other neuron activities σ_j weighted by the couplings J_{ij} , and of the local input h_i .

We compare in Fig. 1b the inferred couplings J_{ij} vs. the logarithms of the correlation indices, $\text{CI}_{ij} = f_{ij}/(f_i f_j)$ for all epochs and sessions. As shown in Methods, in the absence of any indirect (mediated by other recorded cells) effects, J_{ij} and $\log \text{CI}_{ij}$ would be equal. Figure 1b shows that the equality approximately holds for many pairs of neurons with large couplings, either positive or negative. Remarkably, many more pairs of neurons (i, j) have large CI_{ij} (in absolute value) even though $J_{ij} = 0$: for those pairs the apparent correlation (measured by CI_{ij}) is not due to any direct functional connection ($J_{ij} = 0$), but results from the presence of couplings with other third-party cells. The use of couplings rather than correlations offers us a sparser representation of the data (Fig. 1a): about 75% of couplings are inferred to be zero (average fraction over all epochs and sessions), while about 40% of pairwise correlations are compatible with zero within one standard deviation.

Once the coupling and local input parameters are inferred, we may sample the model distribution P through Monte Carlo simulations to check how the statistics of the data is reproduced by the model. The quality of the reproduction of the single-neuron and pairwise spiking probabilities in a time bin is shown in Fig. 2a for the Task epoch of one particular session, which we call A. We can then use P to make predictions for higher-order moments, such as triplet firing probabilities and the probability of multiple-neuron firing in a time bin. Results are compared to the same quantities computed from the spiking data in Fig. 2b. The quality of the inferred distribution P is then assessed through a 4-fold cross-validation procedure: we divide the data set into a train set (3/4 of the time bins) and a test set (1/4 of the bins). The good agreement between the values of observables in the train and test sets in Figs. 2a&b confirms the absence of overfitting in our inference. We show in Fig. 2c the probabilities of the 2^{10} configurations of firing of one subset of ten cells. The Ising model predictions are in much better agreement with the data than the independent-cell model, which reproduces the single-neuron spiking probabilities f_i only. Taking into account pairwise correlations through the effective couplings J_{ij} is therefore crucial to better fit the neural activity distribution.

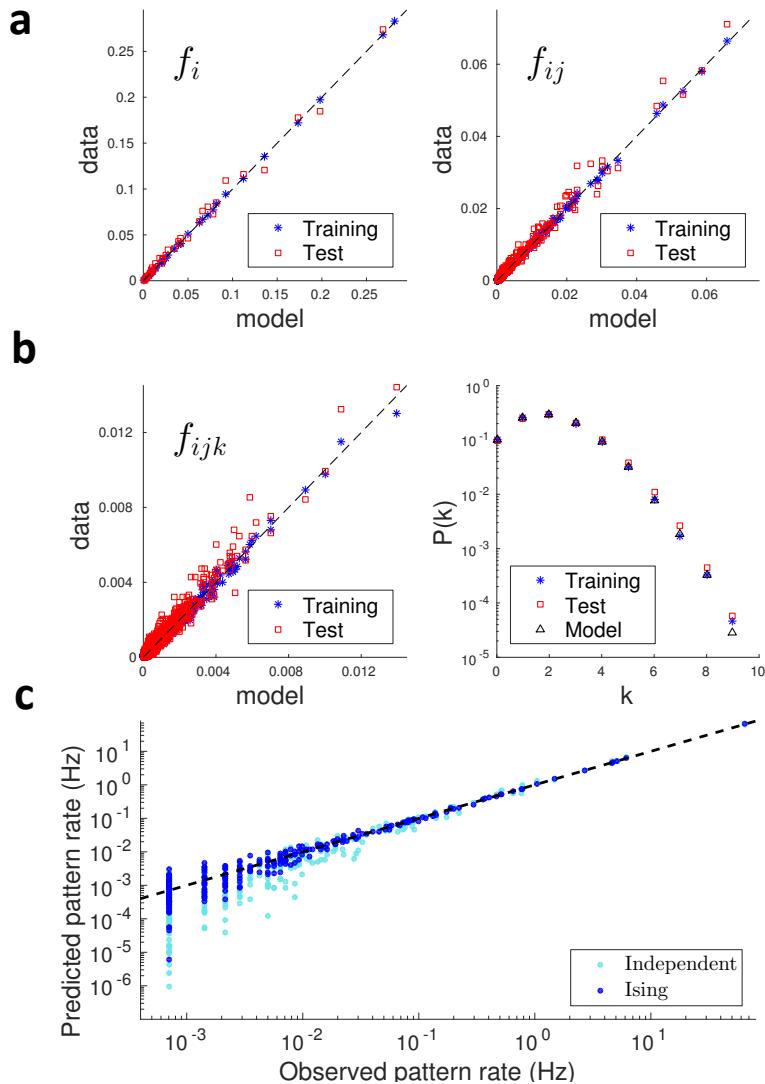


Figure 2. Quality and validation of the inferred model. Reproduction of the statistics of the spiking data for the Task epoch of session A. All panels compare the values of the observables computed from the spiking data with their counterparts computed from the inferred model distribution P , Eqn. [1]. **a.** firing probabilities of single (f_i , left panel) and pairs of (f_{ij} , right panel) neurons. The agreement between the spiking probabilities computed from the data and from the inferred Ising distribution shows that the inference procedure is accurate. 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). **b.** probabilities of firing for triplets (f_{ijk} , left panel) of neurons, and of k neurons to be simultaneously active in a time bin of duration $\Delta t = 10$ ms (right panel). The agreement between the data and model multiple-neuron firing probabilities ($p(k)$) is very good as long as $p(k)$ times the number of time bins in the recording is > 1 , that is, provided the recording time is sufficient to sample rare configurations of multiple neuron firing. Same 4-fold cross-validation procedure and symbols as in Fig. 2a. **c.** probabilities of the $2^{10} = 1024$ activity configurations over a subset of ten cells in the Task epoch of session A. Blue symbols show the scatter plot for the Ising distribution P (inferred from all recorded data), while cyan symbols correspond to the independent-cell model (with all couplings $J_{ij} = 0$, and local inputs h_i fitted to reproduce the single-neuron probabilities f_i). Similar plots are found for other subsets of ten cells among the $N = 37$ recorded cells.

Comparison of functional couplings across epochs shows learning-related potentiation.

The distributions (over all sessions) of inferred coupling parameters are similar across epochs, see Fig. 3a. In addition, little variation over the magnitudes of couplings is observed from session to session. As an illustration we report in Fig. 3a the histograms of coupling parameters for session A. Due to the smaller number of data the histograms are less smooth than the average distribution over all sessions, but span the same ranges of values for J . Similar results hold for local inputs, see SI, Fig. S2.

Despite the overall similarities between the coupling distributions across epochs subtle patterns can be observed when tracking the changes in the couplings corresponding to the same pairs of cells across the different epochs of the same session. We partition the set of couplings in each epoch in three classes, according to their values J and statistical uncertainty δJ : couplings reliably inferred as positive, *i.e.* such that $J/\delta J > 3$, define the [+] class, couplings reliably inferred as negative ($J/\delta J < -3$) form the [-] class, and the remaining couplings are gathered into class [0]. Each pair of neurons (i, j) belongs to one of the resulting 27 classes, *e.g.* [- + 0] if J_{ij} is reliably negative in Sleep Pre, positive in Task, and statistically undetermined in Sleep Post.

The fractions of pairs of neural cells in the 27 classes, averaged over all sessions, are shown in Fig. 3b. Due to the sparsity of the inferred functional couplings classes with vanishing couplings, *e.g.* [000] contain most of the cell pairs. We observe the presence of conserved couplings across the three epochs, corresponding to the large fractions of pairs in classes [---] and in [++], compared to a simple null model, in which we pull together all couplings according to their classes in Sleep Pre and Post, irrespectively of the class in Task (rightmost panel in Fig. 3b).

An important feature emerging from Fig. 3b is the presence of task-related effective potentiation in the functional couplings. This effect is visible from the relative enrichment of [0 ++] with respect to the null model (two-sample t-test, $p < 10^{-105}$), while no such enrichment is found for classes [0 - +] and [00+]. In other words, we find that the fraction of pairs of neurons with close-to-zero couplings in Sleep Pre and positive couplings both in Task and Sleep Post is larger than what would be expected from the knowledge of the coupling classes in the Sleep epochs only. Task-related effective depotentiation, corresponding to the enrichment of [0 --], is also found, but with a weaker magnitude ($p < 10^{-8}$).

While the results above were obtained through averaging over all sessions, there are substantial variations in the fractions of pairs in the classes from session to session. We show in Fig. 3c three examples, referred to as sessions A, Y and B. For sessions A and B effective potentiation, represented in particular by class [0 ++], is clearly visible (with, respectively, $p < 1.8 \cdot 10^{-11}$ and $p < 7.4 \cdot 10^{-4}$). Session Y shows a strong effective depotentiation, represented in particular by class [0 --] ($p < 5.9 \cdot 10^{-8}$).

To characterize quantitatively experience-related changes in the functional couplings in each session we introduce the following session-wide effective potentiation, measuring the amount of potentiation in the couplings from Sleep Pre to Sleep Post, coherently with their values in Task:

$$Pot = \sum_{\substack{\text{pairs } i, j \text{ not in [0] classes} \\ \text{in Task and Sleep Post}}} \theta(J_{ij}^{Task} - J_{ij}^{Sleep\ Pre}) \times (J_{ij}^{Sleep\ Post} - J_{ij}^{Sleep\ Pre}) \quad [3]$$

Summation is restricted to pairs i, j of neurons, whose couplings are significantly different from zero both in Task and Sleep Post (same criterion $|J_{ij}|/\delta J_{ij} > 3$ as for the classes above). The presence of the θ function, $\theta(u) = 1$ if the argument $u > 0$ and 0 if $u \leq 0$, restricts contributions to pairs, whose effective couplings increase from Sleep Pre to Task. In practice positive, respectively, negative contributions to Pot come mostly from classes [0 ++], respectively, [0 +-].

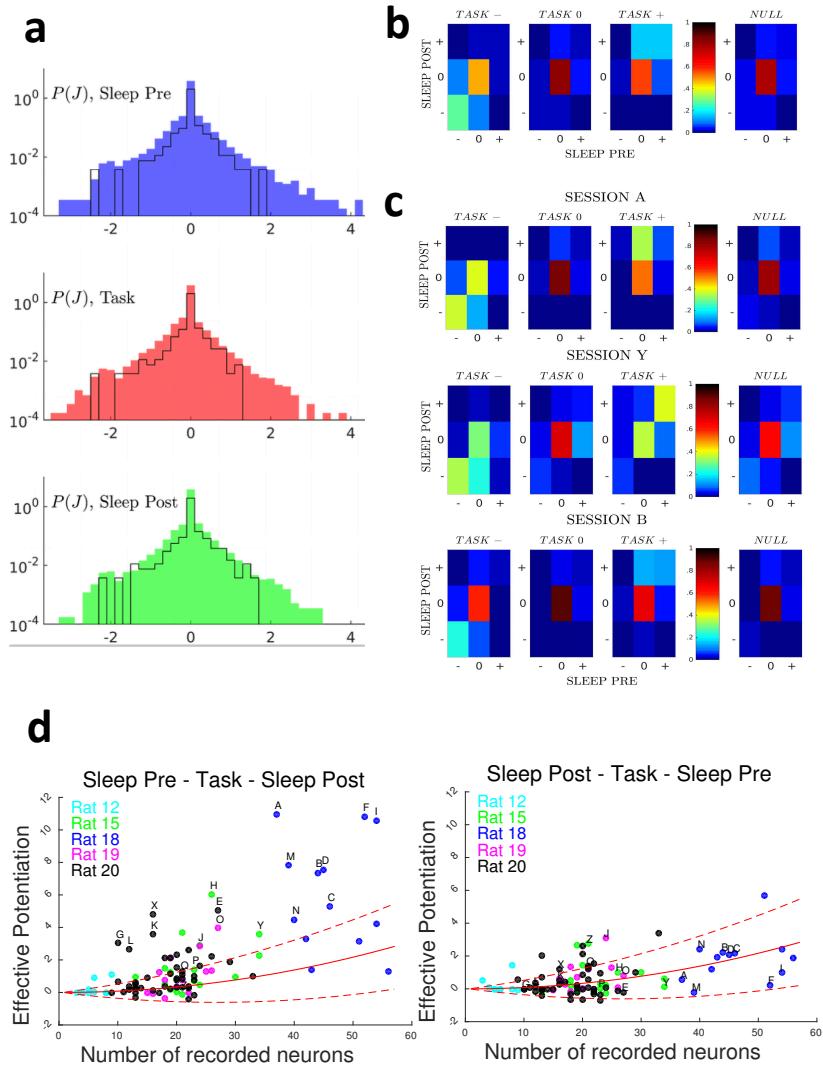


Figure 3. Comparison of couplings across epochs shows task-related effective potentiation.

a. Distributions of the inferred couplings J across the recording sessions for the three epochs (from top to down, Sleep Pre, task and Sleep Post). The black lines show the histograms for Session A only. The sharp peaks in $J = 0$ in the coupling distributions show that the inferred coupling networks are sparse. The average fractions of vanishing couplings are close to 0.75 in all three epochs. The part of the J -distributions corresponding to the minimal value of couplings ($\simeq -2$) permitted by the regularization is due to pairs of cells that never spike together ($f_{ij} = 0$), see Methods. **b.** Fractions of pairs of neurons (i, j) in the 27 classes $[xyz]$, where x, y , and $z = -, 0, +$ are the signs of the inferred couplings J_{ij} in, respectively, Sleep Pre, Task, and Sleep Post. Fractions are computed from all epochs in the 97 recorded sessions. Rightmost column: null model for the fractions corresponding to classes x, z in the Sleep epochs, irrespectively of the y class in Task. **c.** Same as panel b, but for sessions A (top), Y (middle) and B (bottom). **d.** Left: Effective potentiation Pot (Eqn. [3]) is shown for the 97 experimental sessions as a function of the number of recorded neurons identified in all three epochs. Red lines show the predictions of the null model (average: full lines, ± 1 standard deviation: dashed lines), see Methods. Colors identify the five recorded rats; Sessions labelled A to Q (and X) display a large Pot ; sessions Y and Z show depotentiation, see SI, Fig. S11. Right: control case, where we have exchanged the Sleep Pre and Sleep Post inferred couplings.

The effective potentiation Pot is computed for all the 97 recorded experimental sessions, with the results reported in Fig. 3d (left panel). To better assess the statistical significance of the session-wide Pot in Fig. 3d (left) we have built a null model, in which the correspondence between pairs of neurons across the epochs is removed through a reshuffling of the neuron indices and values of couplings are randomly drawn from the distributions of Fig. 3a. Red curves in Fig. 3d show the average value of the effective potentiation within the null model, together with \pm one standard deviation (Methods). For many sessions the effective potentiation is compatible with the expectations of the null model. However, some sessions have large and positive Pot more than one standard deviation above the null model average, see Fig. 3d (left). We label 17 of those sessions (corresponding to $\approx 20\%$ of the recording sessions), with letters A to Q, for further discussion. As a simple control, we have, for each session, swapped the Sleep Pre and Sleep Post couplings in Eqn. [3]. No large- Pot session is found any longer after the swap, see Fig. 3d (right). This simple control provides a clear evidence for the fact that large values of Pot capture experience-related changes in the couplings.

A major source of variability in the data is the limited number of randomly sampled neurons. To assess the influence of sampling on Pot we focus on one particular session (A) with large effective potentiation, and remove cells, one at a time, from the recording. Results are shown in Fig. 4a. In most cases removal of one cell does not affect significantly the value of Pot . A substantial decrease is however observed for a small number of cells, indicated by the labels in Fig. 4a. This results clearly shows that most contributions to Pot come from a restricted subset of the recorded neurons. How many of those relevant cells are well sampled or not may explain, at least in part, the variability in potentiation values observed across sessions.

Groups of neurons supporting effective potentiation are replayed in the sleep epoch after learning

The network of couplings supported by the group of neurons identified in Fig. 4a are shown for the three epochs of session A in Fig. 4b. The effective potentiation from Sleep Pre to Sleep Post, and the strong similarity between the densely interconnected networks in Task and Sleep Post are clearly visible. For this session the potentiated couplings are not supported by independent, non-overlapping pairs of neurons, but are densely interconnecting a restricted group of neurons. We emphasize that experience-related change in the correlational structure of Sleep Post is better seen with effective couplings than with pairwise correlations. For session A again, we show in Fig. 4c (left) the variations of the couplings between the Sleep epochs, $J_{ij}^{Sleep\ Post} - J_{ij}^{Sleep\ Pre}$, vs. their values in the Task epoch, J_{ij}^{Task} . Most contributions to Pot , located in the top right quadrant are supported by the seven cells identified in Fig. 4a (red circles in Fig. 4c). This shows again that the changes experienced by the couplings J_{ij} between the Sleep epochs are positively correlated to their values in Task. Conversely, the same comparison with the CI instead of the couplings J shows a much blurrier picture, see Fig. 4c: The changes in CI between the Sleep epochs do not seem correlated with their values in Task.

While removing one cell at a time is an effective procedure to determine which neurons contribute most to Pot it is computationally demanding. We have therefore developed a fast and fully automatized spectral procedure to directly identify the group of neurons supporting the effective potentiation. Our procedure is based on taking the neurons with largest entries in the top eigenvector of the Pot matrix, whose elements are the contributions to Pot (Eqn. [3]) of the pairs (i, j) , see Methods for details. The top eigenvectors are shown in SI, Fig. S5 for a few sessions. We generally observe a few large entries, and many small ones. We have set a conservative value for the threshold to retain only large entries, see Methods and SI, Section IV for the influence of the threshold value. For session A, our automatized procedure finds the 5 cells (the 7 neurons identified in Fig. 4a, but neurons 11 and 35) that support the couplings contributing most to Pot .

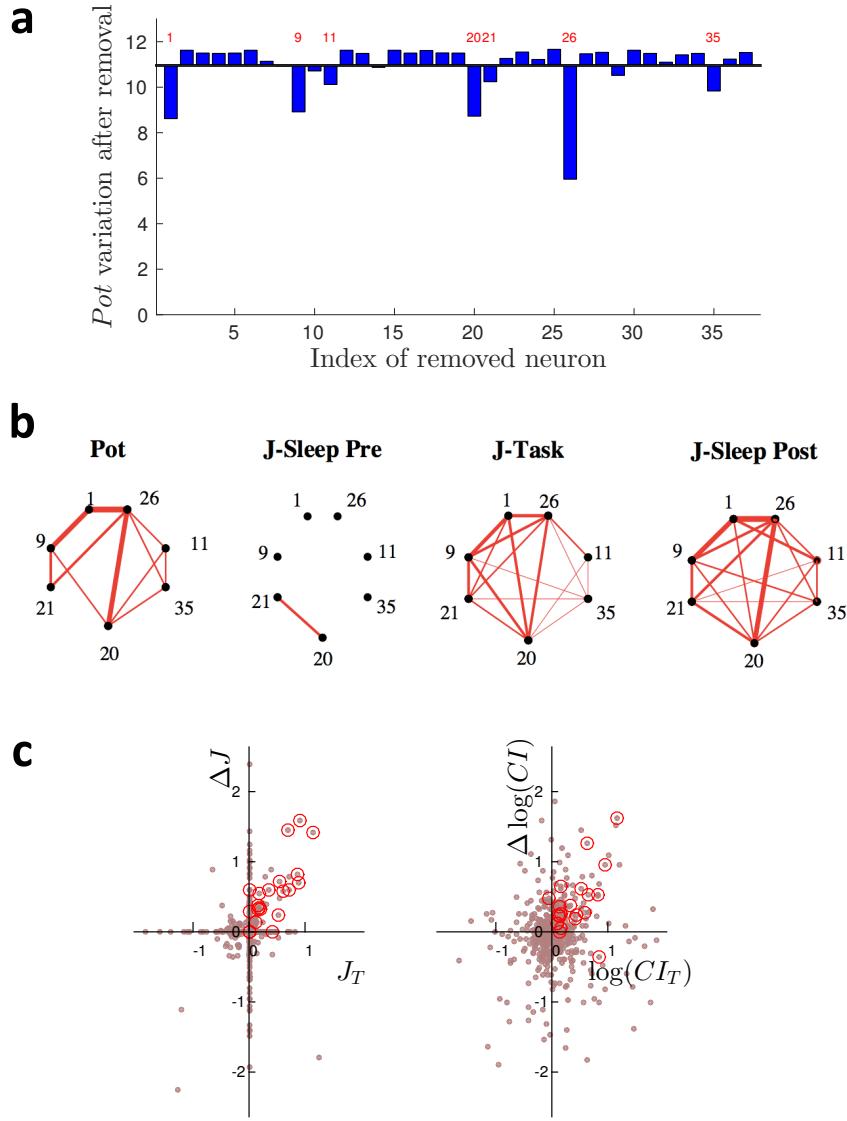


Figure 4. Group of neural cells supporting effectively potentiated couplings in session A.

a. Effective potentiation Pot after removal of one cell (index along the x -axis) from the spiking data of session A. For most cells the removal has no significative impact. A substantial decrease is observed for a few cells, indicated with their indices.

b. Left: main pairwise contributions Pot_{ij} to the effective potentiation (Methods). Right: Networks of couplings supported by the 7-cell group identified in panel a, in the three epochs of session A. Red: $J > 0$, blue: $J < 0$; Line thickness is proportional to $|J|$.

c. Left: Scatter plot of the differences between the couplings in the Sleep epochs, $\Delta J_{ij} = J_{ij}^{Post} - J_{ij}^{Pre}$, vs. their values in Task, J_{ij}^{Task} . Right: same as left panel for the correlation indices CI_{ij} , see text for definition. A group of seven neurons, shown by red circles, supports most of the couplings that are large and positive in Task and in Sleep Post, but not in Sleep Pre. These neurons were identified as the ones giving the largest contribution to Pot , Eqn. [3].

We now show that the groups of neurons identified with our automatized procedure across the 97 sessions really coactivate in the spiking data. To this aim we consider an extension of the pairwise correlation index CI_{ij} to groups of more than two neurons. We define the assembly co-activation ratio (CoA) of a group G of neurons over the time scale τ through

$$\text{CoA}(G, \tau) = \frac{f(G)}{\prod_{i \in G} f_i}, \quad [4]$$

where $f(G)$ is the probability that all the neurons in the group are active within the time scale τ , and the denominator is the product of the individual spiking probabilities. For a group of independent cells the CoA is on average equal to unity. CoA is a very stringent measure of coactivation, as it takes into account only events in which *all* the neurons in the potentiated group are active. A less restrictive measure of the activity of the potentiated group will be studied in the next section.

We first show in Fig. 5a the CoA of the 5-cell potentiated group of session A above. This 5-cell group is found to strongly coactivate in Task on a $\tau \simeq 20 - 40$ ms time scale, and in Sleep Post on a similar time scale, $\tau \simeq 30 - 50$ ms. The 5-cell group does not coactivate more than expected by chance in Sleep Pre, which is compatible with the independent-cell hypothesis due to the low firing frequencies (Methods). This result shows that the potentiated group is replayed in Sleep Post. Interestingly, the coactivation of the potentiated group in Sleep Post is much stronger during non-REM-Sleep periods (non-REM), in which hippocampal sharp waves are known to be important for memory consolidation (Fig. 5a, right). In addition the large CoAs of the potentiated group found in Task and Sleep Post are significantly higher than CoAs for random groups of five neurons (SI, Fig. S10). Those findings suggests that the 5-cell group is (part of) a cell assembly that is reinforced by experience.

For each recorded session we then measure the maximal values (over the time scale τ) reached by the CoA of the group supporting the effectively potentiated couplings in the Sleep Pre and in Sleep Post epochs. The ratio of the two CoAs is a measure of the reinforcement of the coactivation between the neurons in the group across the two sleep epochs in a session. Figure 5b shows the scatter plot of the log. ratios of those two CoAs vs. the values of the effective potentiations Pot of the groups (defined as the sums of contributions to Pot over the pairs of neurons in the groups) across the recorded sessions. A clear monotonic trend is observed, showing that our estimate of coupling potentiation is a good estimator of the existence of neural groups in the spiking data, which reinforce their coactivation in the sleep epoch following task-learning. The sizes of the groups in sessions A-Q, with large overall potentiation in Fig. 3d, range from 2 to 7 cells. It is important to notice that variants of the identified potentiated groups with, say, one more or less cells, can also have large CoAs. Varying the threshold used in the clustering procedure allows us to explore these alternative groups in each session. Examples are provided in SI, Figs. S9-S10.

While we have focused on effective potentiation so far, effective depotentiation, despite being weak on average (Fig. 3b) may be found in some sessions, such as Y, see Fig. 3c. An analysis of the $[0 - -]$ coupling class in session Y has permitted us to identify a cluster of 3 cells. We show in Fig. 5a that this 3-cell group is associated with a decrease of the CoA from Sleep Pre to Sleep Post in non-REM and at short time scales $\tau \sim 20 - 40$ ms. Equation [3] for Pot can be straightforwardly modified to define the effective depotentiation, see SI, Section IV and Fig. S11. In addition to Y we have identified another strongly depotentiated session, Z, and two sessions, C and I, showing both potentiation and depotentiation across their coupling networks. As it is statistically hard to reliably estimate low CoA values, a systematic study of depotentiation across all sessions is difficult, and would require longer recordings.

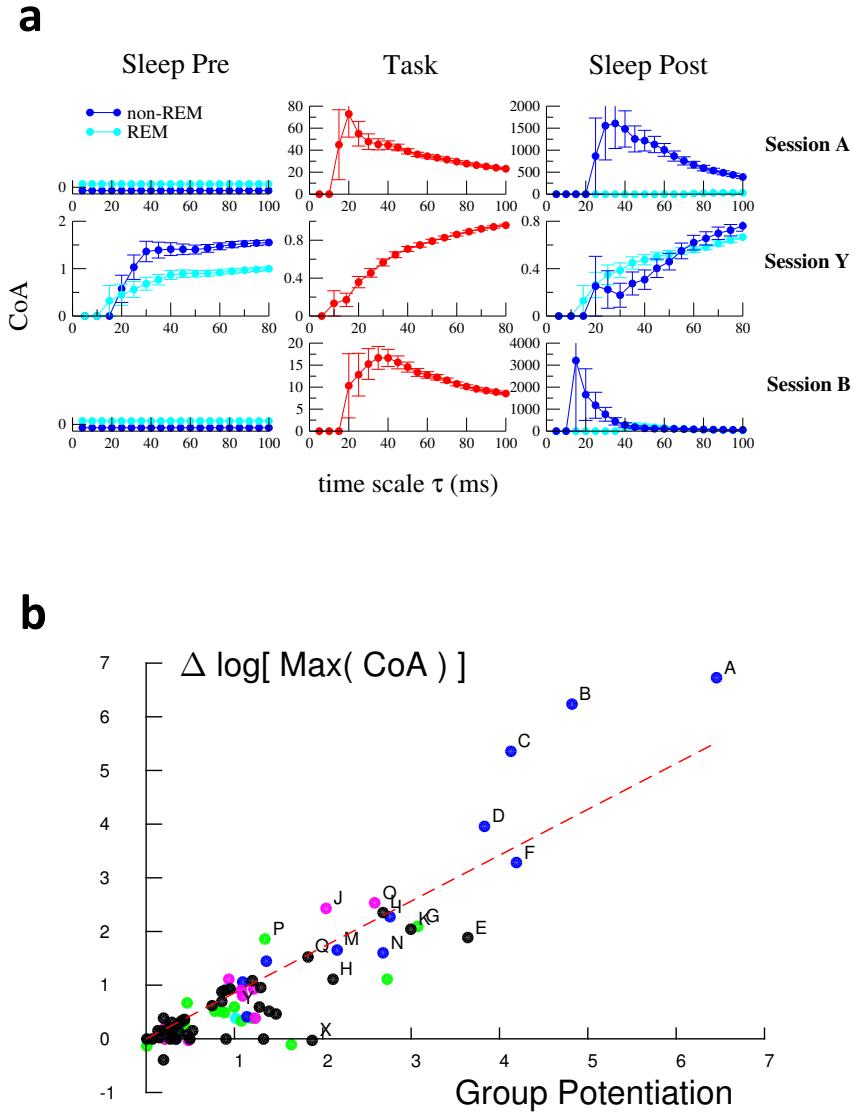


Figure 5. Neurons supporting effectively potentiated couplings show strong increase in coactivation. **a.** Assembly Coactivation ratio (CoA) for groups supporting the effectively potentiated networks of session A, Y, and B. Light and dark blue curves show the CoA for, respectively the REM and non-REM periods of the Sleep epochs. CoAs are shown for time scales τ ranging from 5 ms to $n \times 20$ ms, where n is the number of neurons in each group considered. Note the variations in the CoA and temporal scales along the y - and x -axis between the panels. See Methods for the computation of error bars. CoA equal to zero are compatible with the independent-cell hypothesis due to the low-firing rates of the neurons, see Methods. The potentiated groups of sessions A, B and Y include, respectively, 5, 5 and 4neurons. **b.** Logarithmic change in the peak CoA of the identified potentiated group between Sleep Pre and Sleep Post vs. potentiation Pot for all experimental sessions with at least 2 cells in the potentiated group. The sessions labelled from A-Q have a large group potentiation and a large logarithmic change in their peak CoA ($\Delta[\log \text{Max}(\text{CoA})] > 1.1$). The straight line shows a linear fit of the results (slope = 0.85, $R^2 = 0.8$, p-value = 10^{-32}), proving that large Pot correspond to strong differences in coactivation in the spiking data between the Sleep Pre and the Sleep Post epochs.

Dynamics of reactivation: Effects of hippocampal ripples and connection with learning

The previous analysis allowed us to identify effective potentiated groups that are strongly coactivated in Sleep Post in 17 sessions (A-Q). A fundamental issue is whether the reactivation of those potentiated groups is mostly triggered by hippocampal inputs (sharp-wave ripples, monitored in the experiments, see Methods) or reflects the internal dynamics of the PFC network, modified upon learning. To address this question, for each of the above sessions, we define the Reactivation of the potentiated group in time bin t , $R(t) = \frac{1}{K} \sum_{\ell=1}^K \sigma_{i_\ell}(t)$, where (i_1, i_2, \dots, i_K) are the indices of the K neurons in the group, and the average value of the reactivations over all time bins, $\langle R \rangle$. Reactivation $R(t)$ is less restrictive than CoA in Eqn [4] as it does not require the synchronous co-activation of all the neurons in the group.

We first compute the Ripple-conditioned Reactivation, $RR(\tau)$, defined as the average value of the reactivation following a ripple event by a delay τ , normalized by the average reactivation,

$$RR(\tau) = \frac{1}{N_r \langle R \rangle} \sum_{m=1}^{N_r} R(\tau + t_m) \quad [5]$$

where the t_m s are the times of the N_r ripple events. Figure 6a shows the ripple-conditioned reactivations $RR(\tau)$ for the Sleep Post epochs of sessions A, B, C, which are representative of the variety of RR patterns found across all sessions. In sessions A and B, a marked reactivation peak is found at short time scales of tens of ms. In addition this ‘fast’ peak is followed in session B by a long-lasting reactivation, decaying over a few seconds. No clear response of the reactivation to ripple events is found in Session C on any time scale.

A complementary characterization of the reactivation dynamics, not directly related to ripples is provided by the following normalized Auto-correlation of the Reactivation,

$$AR(\tau) = \frac{\langle R(t) R(t + \tau) \rangle}{\langle R \rangle^2} \quad [6]$$

where the brackets $\langle \cdot \rangle$ denotes the average over all time bins t . AR captures dynamical scales, irrespectively of their origins (inputs from the hippocampus or internal dynamics of the PFC network). The behaviours of $AR(\tau)$ are reported for sessions A, B, C again in Fig. 6a. For all three sessions we observe a large peak in the auto-correlation at $\tau = 0$, expressing the tendency of neurons in the potentiated group to fire together and decaying over a few tens of ms. This ‘fast’ peak is followed by a slow component decaying over few seconds. Remarkably, in session C, for which no reactivation following ripples was detected, AR is stronger than for sessions A and B in spite of a smaller number of neurons in the potentiated group (4 instead of 5 for A and B).

The results above were extended through a systematic analysis of sessions A-Q showing strong effective potentiation. The magnitudes of the fast ($0 < \tau < 50$ ms) and slow ($0.5 \text{ s} < \tau < 1.5 \text{ s}$) components of the ripple-conditioned reactivation are reported in Fig. 6b; results are expressed in terms of Z-scores with respect to a null model defined from the behaviour of $RR(\tau)$ at negative delays $\tau < 0$, see Methods. In Fig. 6c we plot the amplitudes of the fast and slow components to the autocorrelation $AR(\tau)$.

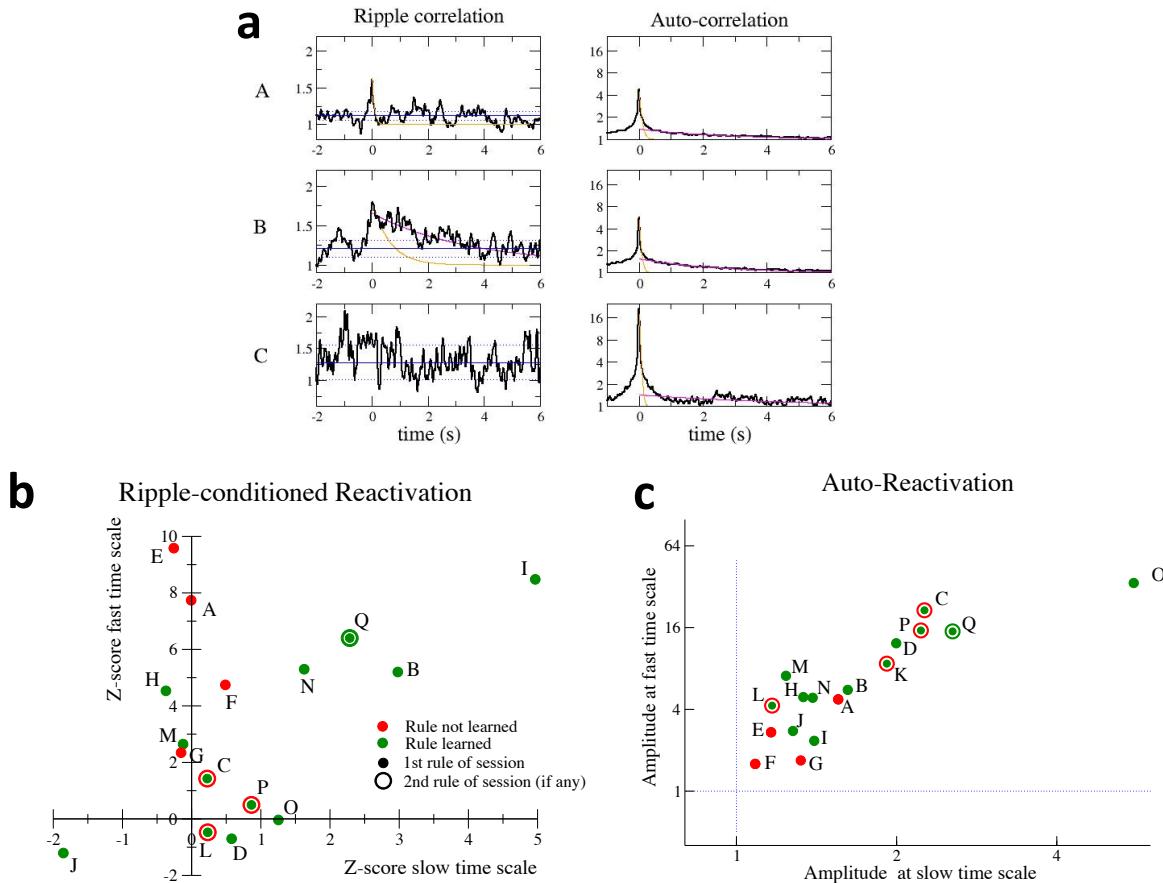


Figure 6. Ripple-Reactivation and Auto-Reactivation of potentiated groups and learning behaviors. **a.** Sliding average over a 50 ms time window of the Ripple-conditioned Reactivation ($RR(\tau)$) in Eqn [5], left panels) and Auto-Reactivation ($AR(\tau)$) in Eqn [6], right panels) of the potentiated groups for sessions A, B, and C. Orange and magenta lines show, respectively, exponential fits of the fast (over the $20 < \tau < 100$ ms range) and slow (over the $1 < \tau < 4$ s range) components to RR and AR . Fast decay times are so fitted to $\tau = 85$ ms and $\tau = 400$ ms for, respectively, sessions A and B, and the slow decay time for session B is $\tau = 3.5$ s. Left: blue lines represent the null model for RR (full line: average value of RR , dotted lines: ± 1 standard deviation), see Methods. Right: blue lines show for comparison the AR curve of the groups of the same size as the potentiated groups and containing the most active neurons in each session. **b.** Z-scores $Z(\tau)$, Eqn [11] in Methods, of the amplitude of the fast ($\tau = 0$) and the slow (averaged over $0.5 < \tau < 1.5$ s) components to the ripple-conditioned reactivation for sessions A to Q. Colors indicate whether the rule was learned (green) or not (red). Sessions in which the rule was changed are represented by two circles (inner circle: first rule, outer circle: last rule), with the same color code. Session K is not shown due to the very small number of detected ripples. Average and standard deviation of the fast decay times over the 8 sessions with top fast response to ripples are $\tau^{RR\text{-fast}} = 230 \pm 180$ ms. Average value and standard deviation of the slow decay time over the 3 sessions (I,B,Q) with a significant slow response to ripples are $\tau^{RR\text{-slow}} = 4 \pm 3$ s. **c.** Amplitudes of the fast (in $\tau = 0$) and the slow (average over $0.5 < \tau < 1.5$ s) components in $AR(\tau)$ for sessions A to Q. Same color code for learning behavior as in panel **b**. Fast and slow decay time constants, fitted over all sessions are, respectively, $\tau^{AR\text{-fast}} = 56 \pm 24$ ms and $\tau^{AR\text{-slow}} = 3 \pm 1.5$ s.

Interestingly the characterization of sessions in terms of *RR* and *AR* seems to be related to the learning behavior of the rat during the Task epoch [27]. The decisions to consider a rule as learned and to change it were based on the success rate of the rat. The rule was said to be learned if the rat did three consecutive correct trials, and had then a success rate larger than 80% over the remaining trials up to the session end or to the change of the rule (learning point criterion). The rule was changed during the session each time the rat did 10 consecutive correct trials, or made only one error over 12 trials (rule-changing criterion), see SI, Section VII, for learning behaviors in sessions A-Q. Different learning scenarios are illustrated by sessions A, B, C studied above (Fig. 6a). In session A the rule was changed at the beginning of the session and was not learned at the end (learning point not reached). In session B the rule was not changed with respect to the previous session, and the learning point was reached by the end of the session. In session C the rule was also the same as in the previous session and was changed in the middle of the session since the rat had fulfilled the rule-changing criterion; the second rule was not learned and the first one was. Despite the limited number of available sessions three empirical observations seem to emerge from Figs. 6b&c:

(a) Sessions in which the learning point was not reached show only a large fast component in the response to ripples (*RR*). Examples are E, A, F, G.

(b) Sessions where the rule has been learned but not changed yet, *i.e.* the success rate has reached intermediary values between the learning and rule-changing points, show also a large slow component in *RR*. Examples are I, B, N, and Q; this observation holds with respect to the second rule for the latter session, see below.

(c) Fast and slow components in the autocorrelation of the reactivation (*AR*) are positively correlated across sessions (Fig. 6c), and have larger amplitudes in sessions where learning has been strong enough to reach both criteria. Examples are C, P, Q, K, D (for the latter session, both criteria were met in the previous session but the rule was not changed in the experiment). Interestingly, over all sessions in which the rule was changed, only the one where the learning point was reached for the second rule (Q) show both a strong *AR* value and a slow component in *RR*.

The complete list of sessions with detailed results can be found in SI, Section VII. Particular cases are further discussed below.

Discussion

In the present work we have focused on experience-related modifications to the functional connectivity matrix (between tens of recorded neurons) in the prefrontal cortex (PFC) of rats [27, 28]. Functional connectivity was here defined through the introduction of a graphical (Ising) model, accounting for the statistical dependencies between spiking events of the neurons in the recorded population. Comparing the functional networks in the two Sleep epochs before and after learning we found, in a substantial fraction of the sessions under investigation, some changes correlated to the functional connectivity during the learning epoch itself. In most of these sessions, couplings became effectively potentiated, and were supported by a limited subset of the neuronal cells (the so-called potentiated group). In other words, a group of cells became much more strongly interconnected in the Sleep epoch after learning than before. We have directly verified on the spiking data that neurons in the identified potentiated groups coactivated much more in the Sleep epoch after than before the learning epoch, which is reminiscent of the notion of cell assembly introduced by Hebb [16] as the basic unit of neural computation and memory. Study of the reactivation dynamics of the potentiated groups allowed us to separate effects due to hippocampal inputs (ripples) or to a putative PFC network, in connection with learning.

Patterns of changes in functional couplings between epochs and potentiated groups

As a general result we have found that functional couplings define sparse interaction network in each single epoch, the class [000] concentrating most of the pairs. In addition there is an overall correlation between the amplitude of couplings across the different epochs, including Task, which can be seen from the relatively large fractions of pairs in classes [−−−] and [+ + +] compared to other non-sparse classes. Classes [0 + +] and, to a lesser extent, [0 − −] contain on average significantly more pairs of neurons than their counterparts [+ + 0] and [− − 0], leading to the general conclusion that some effective couplings undergo substantial Task-related changes from Sleep Pre to Sleep Post. Notice that the classes [+ + +] and [− − −] corresponding to modulations in the amplitudes of the couplings (keeping a fixed sign) across the epochs also contribute to, respectively, effective potentiation and depotentiation. These effects can be analyzed in detail, session by session.

While most of our analysis was focused on experience-related modifications to the functional connectivity, other mechanisms may take place. Tononi and collaborators [35] have suggested that, during specific phases of sleep [37], small synaptic interactions are erased, a phenomenon called homeostasis. The overall similarity between the distribution of inferred couplings in the two Sleep epochs, with many zero couplings, is somewhat in agreement with this hypothesis. However, it is difficult to distinguish between small couplings and couplings strictly equal to zero; homeostatic changes, if any, would likely fall in the most populated [000] class.

The changes in the inferred networks of functional couplings between the Sleep epochs, correlated with the coupling network in the Task epoch, are supported by a subset of the recorded neurons in each session. The identification of the potentiated groups of neurons was then done through an automatized spectral analysis of the *Pot* matrix. We have shown that the groups of potentiated neurons strongly coactivate in the Sleep epoch posterior to learning, and are therefore part of a replayed experience-related cell assembly. It is clear however that the notion of potentiated group should be intended in a statistical sense. Slight variations in the composition of the group, such as adding or removing one specific neuron are associated to large coactivation, as shown in SI, Figs. S9&S10.

It is a remarkable, and somewhat counterintuitive fact that the network of couplings inferred from pairwise co-activation on short time scales, $\Delta t = 10$ ms, suffices to predict coactivation patterns between n neurons on longer time scales, $\tau \simeq n \times \Delta t$ ms. However, even in case of coactivations events, the repeated spiking of neurons in short ‘bursts’ generates a sequence of pairwise coactivation events (Fig. 1a), and the coactivated groups appear as strongly interconnected . Robustness of predictions against the global temporal scale of the cell assembly and the activation ordering is an important advantage of the Ising model, because cell assembly can be played and replayed at different time scales (and in direct and reverse orderings).

Functional couplings: consequence of common inputs or real interactions?

As first discussed in the works of Gerstein and collaborators [10,11] functional couplings can either reflect synaptic connections or the presence of a transient common input co-activating two or more neurons. Within the Hebbian paradigm, coactivation is a prerequisite to learning, favoring synaptic potentiation, *e.g.* through LTP [15]. In our data common inputs could be identified in the transient sharp waves from hippocampus to the prefrontal cortex, occurring preferentially during nonREM Sleep. Sharp-wave ripples have been experimentally demonstrated to be essential for memory consolidation [15,36,37]. Synaptic potentiation in the cortex has been suggested to take place in the immediately following stage, thanks to spindle oscillations contributing to the shutdown of the transmission from the hippocampus to prefrontal cortex. The calcium influx taking place during spindle oscillation could facilitate synaptic potentiation between cells in the replayed assemblies [37,38].

An important issue is whether changes in the functional couplings between the Sleep epochs reflect

such common inputs, necessary for learning, or ‘real’ plasticity in the synaptic interactions. In their earlier work Peyrache and collaborators estimated, based on principal-component analysis, the average reactivation over all sessions, and showed that it occurred within a 2-second time window centered around the sharp-wave event (see Fig 5b in [27]). In our study we computed the reactivation of the more precise potentiated groups defined by the Ising model, for each one of the 17 sessions A-Q. In sessions with a clear response to ripple this response starts within 250 ms from the ripple event, and thus shows a finer temporal resolution. Moreover different scenarios seem to emerge, depending on the learning stage:

For the four sessions in which the rat has not learned the rule the ripple-conditioned reactivation $RR(\tau)$ of the identified cell assembly decays after a short delay of the order of $\tau \sim 200$ ms, comparable to the typical duration of sharp-wave ripples. Hence, the strongly interconnected effective network we identified in Sleep Post (Fig. 4d) seemingly mostly accounts for correlations produced by neural coactivation under common hippocampal inputs.

Furthermore, in several sessions with strong effective potentiation and in which the rule (or two rules in session Q) is (are) learned towards the session end hippocampal ripples evoke a persistent reactivation, lasting several seconds after the ripple event. This effect may signal the existence of an established synaptic potentiation, able to reverberate the activity seconds after the input is over. In other sessions in which the rule was definitively learned and changed no significant reactivation of the potentiated group following the ripple events was found; however, reactivation in those sessions showed a large auto-correlation, decaying over seconds. This observation is apparently contradicted by sessions J and O. However, a thorough investigation of the trial-by-trial learning behavior (SI, Section VII) shows that the binary classification (rule learned/not learned) above may be inadequate. J and O were recorded on the third and fourth days of the same rule, in the same rat (see Fig. 5b). 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 consolidates the learning on day 4 (O) (8 correct trials at the end) but does not reach the rule-changing criterion. It is therefore not surprising that O shows up among sessions with consolidated learning in Fig. 6c.

A tentative interpretation of these findings is the following. For sessions in which the rule has not yet been learned, the large co-activation of the experience-related group, as evidenced in Fig. 5a, seems to be largely supported by the inputs coming from the hippocampus during the sharp wave ripples, known to be crucial for memory consolidation. Conversely, our finding suggests that when the rule has been learned, reactivation occurs over long time scales with two possible mechanisms: slowly decaying persistence of ripple-induced activity, which is found in sessions where the rule has ‘just’ been learned, or second-long replay periods, unrelated to ripple events, which takes mostly place in sessions where the rule has been learned and changed (more stringent criterion). The presence of long-lasting reactivation suggests the existence of a potentiated synaptic network connecting the PFC neurons. This putative network could be either evoked by ripples or subject to spontaneous excitations, depending on its maturity. A possible interpretation of the absence of ripple-induced activity for sessions in which the rule has been definitively learned is that memory has been consolidated, so ripples are not needed for the passage from short-to long-term memory any longer.

The above results are also consistent with the finding that, when the rule has not been learned the CoA of the potentiated group is often substantially larger in non-REM than in REM, see for example CoA of session A in Fig. 5a. In sessions in which the rule has been learned, on the contrary, there is a still smaller but significative coactivation of the potentiated group also in REM periods, as happens in Session B on the time scale of $\tau \sim 50$ ms (see also Fig. S8 for other sessions). Ripple events are indeed more frequent in non-REM sleep. Finally the presence of a CoA larger than one in Sleep Pre, which further increases in Sleep Post, is also often present when the rule is not new (see Fig. S8). New experiments and a larger statistics of data for the different learning and rule shift protocols, would be important to confirm these findings.

Note that, in a related work [39], we have shown how the simulation of the inferred Ising model in the

presence of an external stimulation allowed us to reveal groups of coactivating cells, coinciding with the potentiated groups identified here, and to characterize their statistical variations. This external input, introduced in [39] as a mathematical tool to scan rare co-activation events in the Ising model of the population activity, mimicks, in a very crude way, hippocampal inputs to the PFC.

423
424
425
426

Comparison between inference procedures for Ising model

427
428
429
430
431
432
433
434

The couplings defining the graphical model considered in this paper are an extension of the correlation indices first used to quantify functional connectivity [11–13]. Informally speaking couplings can be viewed as a sparse set of ‘direct’ correlations among the population of recorded cells. Even for sessions with few recorded cells the network of couplings is much sparser than its correlation counterpart (Fig. 1a). We have resorted here to the graphical Ising model, which is the Maximum Entropy model reproducing the 1- and 2-cell firing statistics. In this probabilistic framework, different methods exist to infer the couplings parameters.

435
436
437
438
439
440
441
442
443
444

A competitive inference technique is the standard Boltzmann-Machine learning algorithm [40], broadly used in the analysis of retinal multi-electrode recordings [29, 41, 42]. This inference procedure is slow in its naive version but can become efficient with good initial guess of the coupling parameters [32], upon replacement of the gradient descent for the minimization of the cross-entropy with an approximate version of the Newton method [43], or thanks to improvement specific to the sparse activity of neural population [41]. Two other promising methods to fit Ising models from data are the pseudo-likelihood approach [44] and the minimal probability flow [45]. Both approaches use all the data, and not only the first and second moments of the neural activity, to avoid computing the normalization constant Z in the distribution P , see Eqn. [1]. In particular, the minimal probability flow method has been recently applied to multi-layer restricted Boltzmann Machine to model explicitly the different columns in cortical data [46].

445
446
447
448
449
450
451
452
453

Here we have used the Adaptive Cluster Expansion (ACE, Methods and [32,33]), which has been shown to accurately reproduce interaction parameters for synthetic data. The Ising distribution inferred with ACE reproduces also the statistics of retinal or hippocampal recordings of the activity of tens to hundreds of recorded neurons, including high-order moments [32, 47]. More recently this technique was generalized to the case of non-binary but multiple-categorial variable, or Potts model, to model co-evolution in protein sequences [48]. Our inference approach is very fast on these neural data, taking some seconds on a personal computer to infer the input and functional connectivity parameters. It would be possible to use it to identify cell assemblies online. Combined with optogenetics techniques [49] this would open exciting perspectives in the manipulation of cortical cell assemblies in a controlled way.

454
455
456
457
458
459
460
461
462

The Ising model captures the statistics of snapshots of the activities, and, as such, define symmetric functional couplings $J_{ij} = J_{ji}$. It can therefore not be used to study the ordering in the dynamical activation of the neurons. Other functional connectivity-based inference approaches, *e.g.* the Generalized-Linear [50, 51] and Integrate-and-Fire [52, 53] models are designed to infer non-symmetric connectivity matrices from the temporal sequence of spiking events in the neuronal population. In [39], we have inferred the couplings on the same cortical data set with the Generalized-Linear Model, and found that they are essentially symmetric, and strongly correlated with their Ising counterparts. One possible explanation is that cell assemblies in the prefrontal cortex may also code for intrinsically non-temporal aspects of the task to be learned, in agreement with the findings of [28].

454
455
456
457
458
459
460
461
462

Comparison with existing procedures to identify cell assemblies

463
464
465
466
467
468

Many of the currently available methods to detect and characterize the replay of neural groups, or cell assemblies rely on the knowledge of how the neural activity correlates with sensory or internal inputs. For example, place cells in the hippocampus are known to encode location in space, and replay of place-cell assemblies representing behaviourally meaningful trajectories can be determined with template-matching techniques. More precisely, the ordered activation sequences of place cells observed during salient moments,

such as Sharp-Wave Ripple events, during sleep or wakefulness, is matched with the sequences of place cells determined by the templates observed during locomotion [20,23,54–57]. Similarly, in sensory systems cell assemblies can be detected and characterized by studying the neuronal population response to specific stimuli, easily reproducible in experimental settings. An example is provided by the analysis of neural activity patterns following specific sounds in the auditory cortex [58]. However, those approaches are not easily applicable to all the regions of the brain. In the prefrontal cortex, for instance, neurons may not be activated in a well-defined temporal order, predictable from the knowledge of external stimuli. Cell assemblies might respond to internal cognitive states, or to a combination of extrinsic covariates and internal states, which are very difficult to determine and control experimentally.

In this context, Principal Component Analysis (PCA) has been used as a way to build approximate templates from the correlational structure of data (top principal components), and to detect reactivation, or replay, of those templates. Though PCA was applied successfully to detect replay [27,28,59] it lacks any probabilistic framework and the interpretation of the large entries of the top components is difficult, even with the use of clustering procedures, such as the Assembly Vector estimation [25]. Our analysis significantly extends the principal component analysis (PCA) of [27,28], as it identifies the neurons participating to replay-related assemblies in a detailed way. Let us stress that the whole approach for computing functional connectivity and identifying cell assemblies is fully automatized, and requires spiking data only. While our approach identifies a single cell assembly that contributes the most to functional coupling potentiation, it could be easily extended to the case of more assemblies, see Methods and [61]. Such an extension could be useful for analyzing bigger recordings in future.

A community detection technique for cell assembly identification, exploiting the Markov Stability method, was recently introduced in [60]. The method consists in finding a stable partition on a correlation graph and was tested on hippocampal and retinal data. This graph, unlike the coupling networks we infer here, is defined heuristically and does not disentangle direct (giving rise to coupling) from indirect (mediated through other neurons) correlations.

Materials and Methods

Description of experiments.

Experimental methods were described in details in [27,59] and are summarized in the following. Four Long-Evans male rats weighing 250–300g at arrival were implanted with tetrodes in the prelimbic (PL) subdivision of the medial prefrontal cortex, and in the intermediate-ventral hippocampus. PL tetrodes were used for recording of single units: signals were band-pass filtered between 600 and 6000 Hz, and spikes were detected whenever the filtered signal exceeded a manually set threshold. The resulting waveform (1.3 ms long) were fed into an automated spike sorting algorithm (KlustaKwik, [62]). Hippocampal tetrodes were only used for Local Field Potentials, for the detection of theta rhythms and sharp waves. non-REM was automatically detected, based on power in the cortical delta band (1–4 Hz), hippocampal theta (5–10 Hz), cortical spindles (10–20 Hz) and speed of head motion, by means of a clustering algorithm.

The rats performed an attentional set shift task on a Y-maze, which is known to require the function of the medial prefrontal cortex (mPFC) in rats [63]. Each recording session consisted of a 20–30 minutes sleep or rest epoch (Sleep Pre epoch) in which the rat remained undisturbed in a padded flowerpot placed on the central platform of the maze, a Task epoch, in which the rat performed the behavioral task described below for 20–40 minutes, and by a second sleep or rest epoch (Sleep Post epoch; same situation as in Sleep Pre) of 20–30 minutes. Rats started each trial in the same arm (the departure arm). One of the two other (choice) arms was illuminated at random (pseudo-random schedule: runs of more than 4 consecutive trials with the same illuminated arm were avoided, as were repeated bouts of imposed alternation between the two arms). After that, the central platform was lowered, allowing the rat to access the choice arms.

Only one of the choice arms was rewarded, according to one of 4 contingency rules. Two contingency

rules were spatially guided (always go to the right arm, or to the left arm), the other two were cue guided (go to the illuminated arm, or to the dark arm). The rule was employed at any given moment in time was not signaled to the rat in any way, so that the animal had to learn the rule by trial and error. Once the rat reached a criterion of 10 consecutive correct trials, or only one error out of 12 trials, the rule was changed with no further cue warning to the rat. Rule changes were extra-dimensional, that is, from a spatially-guided rule to a cue-guided rule, and vice versa. All four rats learned in a consolidated way the Right and Light rules (at least 10 consecutive correct trials), whereas only two learned in a consolidated way the left task and one to go where the light is off.

Inference of Ising model parameters.

Inference procedure. We have inferred the Ising model parameters with the Adaptive Cluster Expansion (ACE) algorithm of [32, 33], available from: <https://github.com/johnbarton/ACE>. ACE computes an approximation for the (cross-)entropy of the Ising model reproducing the data,

$$S_{Ising} = \min_{h,J} \left[- \sum_i h_i f_i - \sum_{i < j} J_{ij} f_{ij} + \log Z(\{h_i, J_{ij}\}) \right], \quad [7]$$

with Z defined in Eqn. [1]. The ACE procedure recursively builds clusters $\Gamma = (i_1, i_2, \dots, i_K)$ of neurons with increasing sizes K , and selects those whose contributions $\Delta S_{Ising}(\Gamma)$ to the cross-entropy exceed a threshold Θ (in absolute value) [32, 34]. The total cross-entropy is approximated by the sum of the cluster entropies over the set of selected clusters, $S_{Ising} \simeq \sum_{\text{selected } \Gamma} \Delta S_{Ising}(\Gamma)$. The optimal value of the threshold, Θ , is chosen so that the Ising model (with parameters realizing the minimum in Eqn. [6]) reproduces the experimental low-order statistics within the expected sampling accuracy. Choosing smaller values for Θ would require more computational efforts, and would overfit the data.

Regularization and statistical error bars. To regularize the minimization problem in Eqn. [7] we (1) remove from the data sets neurons spiking less than ten times in each epoch, *i.e.* such that $f_i < 10/B$, where B is the number of time bins, *i.e.* the duration of the recording divided by the time-bin duration Δt ; (2) add to the right hand side of Eqn. [7] a term $\gamma \sum_{k < l} J_{kl}^2$ penalizing large couplings. In practice we choose

$\gamma = N/(10B)$. This choice ensures that the coupling between pairs of neurons that are never active together in any time bin, and which would, in principle, be assigned an infinitely large negative value, become comparable to the couplings between pairs of neurons with a few common spiking event across the recording.

To quantify the statistical errors on the inferred parameters we evaluate the matrix of the second derivatives of the entropy S_{Ising} with respect to the Ising parameters $\{h_i, J_{kl}\}$, also called Fisher information matrix. The squared statistical error bars $(\delta h_i)^2$ and $(\delta J_{kl})^2$ on, respectively, the inferred local inputs h_i and the couplings J_{kl} are given by the diagonal elements of the inverse matrix of the Fisher information matrix, divided by the number of time bins, B [34]. Remark that, with the addition of the regularization term, the Fisher information matrix is definite positive, and, therefore, can be inverted.

Relationship between couplings and correlation indices. In the case of $N = 2$ recorded neurons only the normalization coefficient Z in Eqs. [1] and [7] reads $Z[h_1, h_2, J_{12}] = 1 + e^{h_1} + e^{h_2} + e^{J_{12} + h_1 + h_2}$; the coupling between the two cells is easily obtained upon minimization over h_1, h_2, J_{12} in Eqn. [7], and reads

$$J_{12} = \log \left[\frac{f_{12} (1 - f_1 - f_2 + f_{12})}{(f_1 - f_{12})(f_2 - f_{12})} \right] \simeq \log \left[\frac{f_{12}}{f_1 f_2} \right], \quad [8]$$

for time bins sufficiently short, *i.e.* such that $f_{12} \ll f_1, f_2 \ll 1$. Hence the effective coupling coincides with the logarithm of the correlation index CI_{12} [32]. For $N \geq 3$ recorded neurons no simple relationship exists between couplings and correlation indices due to network effects.

Null model for the effective potentiation.

We define the null model for Pot shown in Fig. 3c (red curves). We first introduce the 3-coupling potentiation $p_{ij,kl,mn}$ through

$$p_{ij,kl,mn} = J_{mn}^{\text{Sleep Post}} - J_{kl}^{\text{Sleep Pre}} \quad \text{if } J_{ij}^{\text{Task}} > J_{kl}^{\text{Sleep Pre}}, \quad \frac{|J_{mn}^{\text{Sleep Post}}|}{\delta J_{mn}^{\text{Sleep Post}}} > 3 \quad \text{and} \quad \frac{|J_{ij}^{\text{Task}}|}{\delta J_{ij}^{\text{Task}}} > 3, \quad [9]$$

and $p_{ij,kl,mn} = 0$ otherwise. Note that the effective potentiation in Eqn. [3] is obtained by summing all 3-coupling potentiations $p_{ij,ij,ij}$ with the *same* pairs of neurons in the three epochs of the sessions. In our null model, the coincidence between the three pairs of neurons is removed by picking up the three couplings in Eqn. [9] above uniformly at random among the set of couplings obtained from *all* sessions of *all* rats. The mean and the standard deviation of the reshuffled (mismatched) 3-coupling potentiations p among all the sessions are: $\langle p \rangle \simeq 0.0018$ and $\Delta p \simeq 0.066$. Note that, as the distributions of couplings are similar from one session to another (Fig. 2a), the distribution of potentiations obtained by reshuffling the couplings within each session only, without pooling all sessions together, gives similar results.

For a session with N recorded neurons the null model distribution for Pot is obtained by summing $\frac{1}{2}N(N-1)$ randomly drawn reshuffled potentiations p . The null model average and its standard deviations are thus given by:

$$\langle Pot \rangle = \frac{N(N-1)}{2} \langle p \rangle, \quad \Delta Pot = \sqrt{\frac{N(N-1)}{2}} \Delta p, \quad [10], \quad (1)$$

which define the red lines of Fig. 3c (left). The null models for the controls in Fig. 3c (right) are obtained with the same procedure after swapping the Sleep Pre and Sleep Post epochs in all sessions.

Identification of the group of neurons supporting the most effectively potentiated couplings.

In order to identify the sub-group of neurons that supports the couplings with strongest effective potentiation in a session we consider the contribution to Pot (Eqn. [3]) coming from the pair (i, j) of neurons, $Pot_{ij} \equiv p_{ij,ij,ij}$, cf. Eqn. [9]. This contribution can be seen as an entry of a $N \times N$ -dimensional matrix. This matrix is sparse, square symmetric, and has large entries for neurons (i, j) supporting strongly potentiated couplings. The top eigenvector, $\mathbf{v} = \{v_i\}$, of the matrix is localized over few neurons i , which strongly contribute to Pot (SI, Fig. S5). We define the potentiated group as the set of neurons i with components v_i larger than a threshold value c (ranging between 0 and 1 for normalized \mathbf{v}). In all the experimental sessions considered here this simple spectral analysis gives at most one large and connected neural group. Spectral graph theory offers efficient methods for dealing with more complex data structures on larger data sets, including more than one largely interconnected group [61].

The correlation between the changes in log CoA from Sleep Pre to Post, $\Delta \log \text{CoA}$, and the potentiation of the group across all sessions (shown in Fig. 5b for $c = 0.22$) varies with c . The best correlations are found in practice in the range $0.15 < c < 0.35$, with p-values ranging between 10^{-30} and 10^{-35} (SI, Fig. S6). We have arbitrarily set $c = 0.22$ in between these two limits. This value is also adequate if one imposes in addition that potentiated groups should include 3 cells at least (SI, Fig. S6).

Statistical significance of the Coactivation ratio (CoA).

To assess the statistical validity of the CoA defined in Eqn. [4] for a group G of neurons we compute the error bar on CoA, shown in Fig. 5a. Assuming a Poisson distribution for the co-activation events, the standard deviation of the CoA is estimated to be $\text{CoA}(\tau)/\sqrt{N_G(\tau)}$, where $N_G(\tau)$ is the number of coactivation events for the cells in G over the time scale τ .

Note that simultaneous-firing events (contributing to $f(G)$) are unlikely to be found, and the CoA is likely to be zero, if the duration of the recording is small, *e.g.* compared to $T_{min} = \tau / \prod_{i \in G} f_i(\tau)$. This happens for the 5-cell potentiated group of session A for time scales $\tau \leq 40$ ms in the Task epoch, and for all the values of τ considered in Sleep Pre and Post in Fig. 5a.

Analysis of ripple-conditioned reactivation.

We define a null model for the average response to ripples after a delay τ , $RR(\tau)$ in Eqn [5], as follows. For each session, we compute the average value, RR_0 , and the standard deviation, δRR_0 , of $RR(\tau)$ over the -3 s $< \tau < -0.5$ s range. The range of delays is sufficiently negative to exclude any inaccuracy in the determination of the ripple times t_m . The values of RR_0 , and $RR_0 \pm \delta RR_0$ are shown in Fig. 6a (left panels) for sessions A , B and C. The Z-score of ripple-conditioned reactivation for positive delay τ is defined through

$$Z(\tau) = \frac{RR(\tau) - RR_0}{\delta RR_0} . \quad [11]$$

The value of the Z-score in $\tau = 0$ and its average over the 0.5 s $< \tau < 1.5$ s interval are used to estimate the amplitudes of the, respectively, fast and slow components in RR , see Fig. 6b.

Supporting Information Files

Text S1. Supporting Information Text for Functional Coupling Networks Inferred from Prefrontal Cortex Activity Show Learning-Related Effective Plasticity 605
606
607

Tables S1. Supporting Information Tables for Functional Coupling Networks Inferred from Prefrontal Cortex Activity Show Learning-Related Effective Plasticity 608
609

Acknowledgments

We thank Georges Debregeas for useful comments about the manuscript.

611

References

1. Friston, K.J. Functional and Effective Connectivity: A Review. *Brain Connectivity* **1**, 13 (2011).
2. McNaughton, B.L., O'Keefe, J. and Barnes, C.A. The stereotrode: a new technique for simultaneous isolation of several single units in the central nervous system from multiple unit records. *J Neurosci. Methods* **8**: 391-397 (1983).
3. Meister, M., J. Pine, and D. A. Baylor Multineuronal signals from the retina: acquisition and analysis. *J. Neurosci. Methods* **51**, 95-106 (1994).
4. Wagenaar, D. A., Pine, J., and Potter, S. M. Effective parameters for stimulation of dissociated cultures using multi-electrode arrays. *J. Neurosci.. Methods* **138**, 27-37 (2004).
5. Wolf, S. et al. Whole-brain functional imaging with two-photon light-sheet microscopy. *Nat. Methods* **12**, 379-380 (2015).
6. Chhetri, R.K., Amat, F., Wan, Y., Höckendorf, B., Lemon, W.C. and Keller, P.J. Whole-animal functional and developmental imaging with isotropic spatial resolution. *Nat. Methods* **12**, 1171-1178 (2015).
7. Nguyen, J.P et al. Whole-brain calcium imaging with cellular resolution in freely behaving *Caenorhabditis elegans*. *Proc. Nat. Acad. Sci* : E1074-E1081 (2015).
8. Seung, H.S. Neuroscience: Towards functional connectomics, *Nature* **471**, 170-2 (2011).
9. Stevenson, I.H., Rebescu, J.M., Miller, L.E. and Koerding, K.P. Inferring functional connections between neurons. *Curr. Opin. Neurobiol.* **18**, 582-8 (2008).
10. Gerstein, G.L. and Perkel D.H. Simultaneously Recorded Trains of Action Potentials: Analysis and Functional Interpretation. *Science* **164**, 828-830 (1969).
11. Aersten, A.M.H.J., Gerstein, G.L., Habib, M.K. and Palm, G. Dynamics of neuronal firing Correlations : Modulation of "effective Connectivity". *Journal of Neurophysiology* **61**, 900-917 (1989).
12. Fujisawa, S., Amarasingham, A., Harrison, M.T. and Buzsáki, G. Behavior-dependent short-term assembly dynamics in the medial prefrontal cortex. *Nat. Neurosci.* **11**, 823-33 (2008).
13. Schwindel, C.D., Ali, K., McNaughton, B.L. and Tatsuno, M. Long-term recordings improve the detection of weak excitatory-excitatory connections in rat prefrontal cortex. *J. Neurosci.* **34**, 5454-67 (2014).
14. Bliss, T. V. P., Collingridge, G. L. A synaptic model of memory: long-term potentiation in the hippocampus. *Nature* **361**, 31-39 (1993)
15. Buzsáki, G., Hippocampal sharp wave-ripple: A cognitive biomarker for episodic memory and planning. *Hippocampus*. **25(10)**:1073-188 (2015) .
16. Hebb, D.O. *The organization of behavior: A neurophysiological approach*, John Wiley & Sons (1949).
17. Harris, K.D., Csicsvari, J., Hirase, H., Dragoi, G. and Buzsáki, G. Organization of cell assemblies in the hippocampus, *Nature* **424**, 552-6 (2003).
18. Buzsáki, G. Neural syntax: cell assemblies, synapsembles, and readers, *Neuron* **68**, 342-85 (2010).
19. Wilson, M.A. and McNaughton, B.L. Reactivation of hippocampal ensemble memories during sleep, *Science* **265**, 676-9 (1994).

20. Carr, M.F., Jadhav, S.P. and Frank, L.M. Hippocampal replay in the awake state: a potential substrate for memory consolidation and retrieval, *Nat. Neurosci.* **14**, 147-53 (2011).
21. O'Neill, J., Pleydell-Bouverie, B., Dupret, D. and Csicsvari, J. Play it again: reactivation of waking experience and memory, *Trends Neurosci.* **33**, 220-9 (2010).
22. Johnson, A. and Redish, A.D. Neural ensembles in CA3 transiently encode paths forward of the animal at a decision point, *J. Neurosci.* **27**, 12176-89 (2007).
23. Pfeiffer, B.E. and Foster, D.J. Hippocampal place-cell sequences depict future paths to remembered goals, *Nature* **497**, 74-9 (2013).
24. Singer, A.C., Carr, M.F., Karlsson, M.P. and Frank, L.M. Hippocampal SWR activity predicts correct decisions during the initial learning of an alternation task, *Neuron* **77**, 1163-73 (2013).
25. Lopes-dos Santos, V., Ribeiro, S.T. and Tort, A.B.L. Detecting cell assemblies in large neuronal populations, *J. Neurosci. Methods* **220**, 149-166 (2013).
26. Billeh, Y.N., Schaub, M.T., Anastassiou, C.A., Baharona, M. and Koch, C. Revealing cell assemblies at multiple levels of granularity, *J. Neurosci. Methods* **236**, 92-106 (2014).
27. 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, *Nat. Neurosci.* **12**, 919-26 (2009).
28. Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S.I. and Battaglia, F.P. Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution, *J. Comput. Neurosci.* **29**, 309-25 (2010).
29. Schneidman, E., Berry, M., Segev, R. and Bialek, W. Weak pairwise correlations imply strongly correlated network states in a population, *Nature* **440**, 1007 (2006).
30. Cocco, S., Leibler, S. and Monasson, R. Neuronal couplings between retinal ganglion cells inferred by efficient inverse statistical physics methods, *Proc. Natl. Acad. Sci. USA* **106**, 14058 (2009).
31. Tkacik, G., Prentice, J.S., Balasubramanian, V. and Schneidman, E. Optimal population coding by noisy spiking neurons, *Proc. Natl. Acad. Sci. USA* **107**, 14419-24 (2010).
32. Barton, J.P., Cocco S. Ising models for neural activity inferred via selective cluster expansion: structural and coding properties *J. Stat. Mech.* **P03002** (2013).
33. Cocco, S. and Monasson, R. Adaptive cluster expansion for inferring Boltzmann machines with noisy data, *Phys. Rev. Lett.* **106**, 090601 (2011).
34. Cocco, S. and Monasson, R. Adaptive cluster expansion for the inverse Ising problem: convergence, algorithm and tests, *J. Stat. Phys.* **147**, 252 (2012).
35. Tononi, G. and Cinelli., C. Sleep function and synaptic homeostasis. *Sleep Med. Rev.* **10**, 49-62 (2006).
36. Girardeau, G., Benchenane, K., Wiener, S.I., Buzsaki, G., and Zugaro, M.B. Selective suppression of hippocampal ripples impairs spatial memory. *Nat Neurosci* **12**: 1222-1223 (2009).
37. Genzel, L., Kroes, M.C.W., Dresler M. and Battaglia, F.B. Light sleep versus slow wave sleep in memory consolidation: a question of global versus local processes? *Trends in Neurosciences* **37**, 1 (2014).

38. Siapas, A.G. and Wilson, M.A. Coordinated Interactions between Hippocampal Ripples and Cortical Spindles during Slow-Wave Sleep. *Neuron* **21**, 1123-1128 (1998).
39. Tavoni, G., Cocco, S. and Monasson R. Neural assemblies revealed by inferred connectivity-based models of prefrontal cortex recordings. *J. Comp. Neurosci.* 1-25, DOI 10.1007/s10827-016-0617-5 (2016).
40. Hinton, G.E. and Sejnowski, T.J. Learning and relearning in Boltzmann machines, *Neural Computation* **1**, 282-317 (1986).
41. Broderick, T., Dudik, M., Tkacik, G., Schapire, R.E. and Bialek, W. Faster solutions of the inverse pairwise Ising problem. *arXiv.org:* 0712.2437 (2007).
42. Marre, O., El Boustani, S., Fregnac, Y. and Destexhe, A. Prediction of spatiotemporal patterns of neural activity from pairwise correlations, *Phys. Rev. Lett.* **102**, 138101 (2009).
43. Ferrari, U. Learning Maximal Entropy Models from finite size datasets: a fast Data-Driven algorithm allows to sample from the posterior distribution. *Phys. Rev. E* **94**, 023301 (2016)
44. Aurell, E. and Ekeberg, M. Inverse Ising Inference using all the data. *Phys. Rev. Lett.* **108**, 090201 (2012)
45. Sohl-Dickstein, J., Battaglino, P. and DeWeese, M. New Method for Parameter Estimation in Probabilistic Models: Minimum Probability Flow. *Phys. Rev. Lett.* **107**, 11-14 (2011)
46. Kästere, U., Sohl-Dickstein, J., Gray, C.M., Olshausen B.A. Modeling higher-order correlations within cortical microcolumns. *PLoS Comput Biol.* **10**, e1003684 (2014).
47. Posani, L., Cocco, S., Jezek, K. and Monasson R. Pairwise graphical model for brain-state identification: application to spatial representation decoding from hippocampal CA1 and CA3 recordings. preprint (2016).
48. Barton, J.P., Chakraborty A.K. , Cocco, S., Jacquin, H. and Monasson R. On the entropy of protein families. *J. Stat. Phys.* **162**, 1267-1296 (2015).
49. Reutsky-Gefen, I., Golan, L., Farah, N., Schejter, A., Tsur, L., Brosh, I. and Shoham, S. Holographic optogenetic stimulation of patterned neuronal activity for vision restoration, *Nat. Commun.* **4**, 1509 (2013).
50. Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P. and Brown, E.P. A point process framework for relating neural spiking activity to spiking history, neural ensemble, and extrinsic covariate effects, *J. Neurophysiol.* **93**, 1074-89 (2005).
51. Pillow, J., Shlens, J., Paninski, L., Sher, A., Litke, A.M., Chichilnisky, E.J. and Simoncelli, E.P. Spatio-temporal correlations and visual signalling in a complete neuronal population, *Nature* **454**, 995-9 (2008).
52. Koyama, S. and Paninski, L. Efficient computation of the maximum a posteriori path and parameter estimation in integrate-and-fire and more general state-space models, *J. Comput. Neurosci.* **29**, 89-105 (2010).
53. Monasson, R. and Cocco, S. Fast Inference of Interactions in Assemblies of Stochastic Integrate-and-Fire Neurons from Spike Recordings, *J. Comput. Neurosci.* **31**, 199-227 (2011).

54. Brown, E.N., Frank, L.M., Tang, D., Quirk, M.C. and Wilson, M.A. A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells, *J. Neurosci.* **18**, 7411–7425 (1998).
55. Foster, D. J. and Wilson, M.A. Reverse replay of behavioural sequences in hippocampal place cells during the awake state, *Nature* **440**, 680–683 (2006).
56. Diba, K. and Buzsáki, G. Forward and reverse hippocampal place-cell sequences during ripples, *Nature neuroscience* **10**, 1241–1242 (2007).
57. Lee, A. K. and Wilson, M.A. Memory of sequential experience in the hippocampus during slow wave sleep, *Neuron* **36**, 1183–1194 (2002).
58. Bathellier, B., Ushakova, L. and Rumpel, S. Discrete neocortical dynamics predict behavioral categorization of sounds, *Neuron* **76**, 435–449 (2012).
59. Benchenane, K., Peyrache, A., Khamassi, M., Tierney, P.L., Gioanni, Y., Battaglia, F.P. and Wiener, S.I. Coherent theta oscillations and reorganization of spike timing in the hippocampal–prefrontal network upon learning, *Neuron* **66**, 921–936 (2010).
60. Billeh, Y.N., Schaub, M.T., Anastassiou, C.A., Barahona, M. and Koch, C. Revealing cell assemblies at multiple levels of granularity, *Journal of Neuroscience methods* **236**, 92–106 (2014).
61. Fortunato, S., Community detection in graphs. *Physics Reports* **486**, 75–154 (2010).
62. Harris, K. D., Henze, D. A., Csicsvari, J., Hirase, H. and Buzsáki, G. Accuracy of tetrode spike separation as determined by simultaneous intracellular and extracellular measurements. *J. Neurophysiol.* **84**, 401–414 (2000).
63. Birrell, J. M. and Brown, V. J. Medial frontal cortex mediates perceptual attentional set shifting in the rat. *J. Neurosci.* **20**, 4320–4324 (2000).