nature neuroscience

Reactivation of emergent task-related ensembles during slow-wave sleep after neuroprosthetic learning

Tanuj Gulati[1,2], Dhakshin S Ramanathan[1,3,4], Chelsea C Wong[1,2] & Karunesh Ganguly[1,2]

Brain-machine interfaces can allow neural control over assistive devices. They also provide an important platform for studying neural plasticity. Recent studies have suggested that optimal engagement of learning is essential for robust neuroprosthetic control. However, little is known about the neural processes that may consolidate a neuroprosthetic skill. On the basis of the growing body of evidence linking slow-wave activity (SWA) during sleep to consolidation, we examined whether there is 'offline' processing after neuroprosthetic learning. Using a rodent model, we found that, after successful learning, task-related units specifically experienced increased locking and coherency to SWA during sleep. Moreover, spike-spike coherence among these units was substantially enhanced. These changes were not present with poor skill acquisition or after control awake periods, demonstrating the specificity of our observations to learning. Notably, the time spent in SWA predicted the performance gains. Thus, SWA appears to be involved in offline processing after neuroprosthetic learning.

Brain-machine interfaces (BMIs) have the potential to seamlessly merge the computational power of the motor system with that of artificial electronic systems. Research into the development of BMIs has flourished over the past decade, leading to impressive demonstrations of rodents[1–3], non-human primates[4–10] and humans controlling prosthetic devices[11–13] in real-time through modulation of neural signals. This body of work has also identified that adaptive processes, both in the motor system and in the algorithms that transform neural activity into prosthetic control signals (that is, decoders), are essential for achieving stable neuroprosthetic control[3,6,8,9,14,15]. For example, practicing neuroprosthetic control over multiple sessions in the setting of a fixed decoder is known to result in persistent improvements in performance[14]. Moreover, long-term adaptation of decoders at a rate that is seemingly concordant with neural learning rates can be essential for long-term stable performance[15]. Together, these studies underscore the importance of optimal engagement of long-term neural learning mechanisms[9,14,15]. However, the specific neural process that stabilizes a newly acquired prosthetic skill remains incompletely understood.

Given the growing body of literature indicating that cortical SWA is associated with offline processing[16–19], we hypothesized that such processing facilitates direct cortical control of an artificial actuator. Studies in humans have suggested that offline processing during overnight sleep[20–23], as well as during daytime naps[24], facilitates memory consolidation. Motor task performance is also known to improve after sleep versus an equivalent amount of wakefulness[21–25]. Although SWA has been directly implicated in the consolidation of motor skills[26–30], it remains poorly understood how neural ensembles in motor cortex are precisely modified during SWA after skill acquisition. Whereas most theories suggest that reactivation of cortical neurons during SWA is important for consolidation[31], there is, to the best of our knowledge, little experimental support of this for procedural memory formation. In this context, BMIs offer a powerful tool to directly modulate single units, thereby allowing precise characterization of how both task-related (TR) and task-unrelated (TU) cortical units are differentially processed in motor cortex during SWA. In these experiments, rats were trained to control the angular velocity of a feeding tube via direct modulation of cortical units through operant conditioning. We hypothesized that, after neuroprosthetic skill acquisition, TR units would be differentially processed from TU units during SWA.

RESULTS

After implantation of microelectrodes in M1, we trained six rats to exert direct neural control of the angular velocity of a mechanical actuator that could also deliver water (**Fig. 1**). A linear decoder with randomized weights typically converted the firing rates of two randomly selected units (hereafter referred to as direct units) into the angular velocity of the actuator (**Supplementary Fig. 1** and Online Methods). We also recorded multiple other units that were not causally linked to actuator movements (hereafter referred to as indirect units). The decoder weights were held constant during the session to exclusively rely on neural learning mechanisms. Each trial started with the simultaneous delivery of an auditory tone and the opening of a door to allow access to the tube (**Fig. 1a**,**b**). At the start of each trial, the angular position of the tube was set to 0° (P_1). If the angular position of the tube was held for >300 ms at position P_2 (90°), a defined amount of water was delivered (that is, successful trial). A trial was stopped if this was not achieved within 15 s (that is, unsuccessful trial). At the end of a trial, the door was closed and the actuator was returned to position P_1.

Over the course of a typical 2-h practice session, animals showed improvements in task performance, with a significant reduction in the time to successful trial completion and a decrease in the number of unsuccessful trials (**Fig. 1c**,**e**). Specifically, of 19 such training

[1]Neurology and Rehabilitation Department, San Francisco VA Medical Center, San Francisco, California, USA. [2]Department of Neurology, University of California, San Francisco, California, USA. [3]Department of Psychiatry, San Francisco VA Medical Center, San Francisco, California, USA. [4]Department of Psychiatry, University of California, San Francisco, California, USA. Correspondence should be addressed to K.G. (karunesh.ganguly@ucsf.edu).

Received 8 April; accepted 12 June; published online 6 July 2014; doi:10.1038/nn.3759

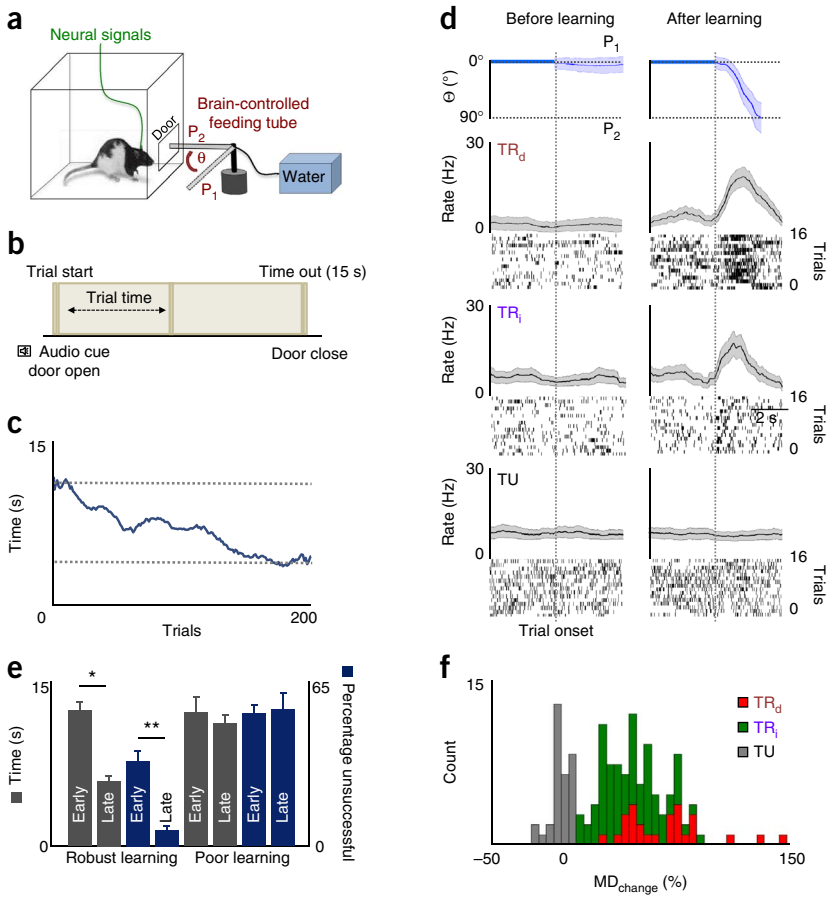

Figure 1 Direct control of motor cortex units. (**a**) Direct neural control of a feeding tube (θ = angular position). Each trial started with the tube at P_1. (**b**) Trial started with an audio tone cue and opening of the door. A successful trial required movement of the tube to P_2 within 15 s. (**c**) Change in task completion time as a function of trial number (line shows moving average of 20 trials). (**d**) Angular position of the tube is shown from a single session. Peri-event histogram from early and late trials from a single session are shown in left and right panels, respectively. Thick line represents mean and shaded area is s.e.m. (**e**) Comparison of time to trial completion and change in percentage of unsuccessful trials for robust learning and poor learning sessions (mean ± s.e.m.). In robust learning sessions, time to reward reduced significantly over trials (*t* test, $t_{14} = 7.31$, $^{*}P < 10^{-5}$) and the proportion of unsuccessful trials also dropped significantly (*t* test, $t_{14} = 8.58$, $^{**}P < 10^{-6}$). In poor learning sessions, time to reward did not change significantly (*t* test, $t_3 = 1.73$, $P = 0.18$). Similarly, the proportion of unsuccessful trials did not change significantly (*t* test, $t_3 = -0.52$, $P = 0.68$). (**f**) Histogram of percentage change in modulation depth (MD_{change}) for each of the three categories of units.

sessions, there were 15 'robust learning' sessions with significant reductions for both metrics (n = 15 sessions; early completion time, 12.42 ± 0.74 s; late completion time, 5.93 ± 0.48 (mean ± s.e.m.); *t* test, $t_{14} = 7.3$, $P < 10^{-5}$; **Fig. 1e**). Moreover, the percentage of unsuccessful trials significantly decreased from 33.56 ± 4.04% to 6.00 ± 1.51% (*t* test, $t_{14} = 8.5$, $P < 10^{-6}$). For the remaining four sessions, these values did not improve with training (that is, poor learning). In the robust learning sessions, we also found that the path of the actuator taken from position P_1 to P_2 became direct (**Fig. 1d**).

For direct units that were causally linked to actuator movements via the decoder (that is, TR direct or TR_d), nearly all experienced significant changes in TR modulation at the end of session (96% of 27 TR_d units with a significant increase in modulation depth). We assigned weights in the decoder that ranged from +1 to −1. Although we did not see suppression of activity relative to the baseline, there was a significantly greater modulation of the units assigned positive weights (that is, TR_d^+) than negative weights (that is, TR_d^-) (*t* test, $t_{25} = 5.17$, $P < 10^{-4}$; **Supplementary Fig. 2a–c**). Direct units were also significantly more likely to experience a change in modulation than indirect units (n = 15 sessions; direct, 96.7% ± 3.5%; indirect, 65.7 ± 3.7%; *t* test, $t_{28} = 5.84$, $P < 10^{-4}$). However, given the growing notion that subsets of indirect units likely contribute to neuroprosthetic control[32–34], we further subclassified such units as being either task-related (TR_i) or TU based on changes in modulation depth with learning. A unit was declared TR_i if its post-learning change in firing rate was 2.5 s.d. above the baseline firing rate (Online Methods). Moreover, we did not find a preponderance of particular cell types in either category, as determined by the width of the record spike (**Supplementary Fig. 1b–d**). For subsequent analysis, we respectively used 27 TR_d, 108 TR_i and 39 TU units in 15 BMI training sessions (**Supplementary Table 1**). TR (TR_d and TR_i) units experienced large changes in modulation depth (**Fig. 1f**). Notably, after learning, the temporal profiles of activation (**Fig. 1d**) for TR_d and TR_i were not significantly different (time to peak firing: direct, 3.1 ± 0.2 s; TR_i, 3.3 ± 0.1 s; *t* test, $t_{133} = 1.18$, $P = 0.24$), suggesting that these two distinct groups of TR units became functionally and temporally coupled with learning and successful task performance. We also saw evidence for changes in the spike-triggered local field potential (LFP) for TR units that was not present for the TU units (**Supplementary Fig. 2e,f**).

Increased locking to SWA after learning

Given that previous studies have suggested a link between slow-wave sleep (SWS) and motor learning[16–18], we first examined whether the spike-field relationship during SWS is altered after robust learning sessions. To look for signatures of offline processing, we first assessed whether individual units (TR_d, TR_i and TU) experienced changes when comparing the pre- and post-SWS around a successful training session (hereafter referred to as SWS_{pre} and SWS_{post}; Online Methods and **Supplementary Fig. 3**). Specifically, we used spike-triggered averaging (STA) to quantify the relationship between spiking and SWA. The STA provides an intuitive estimate of how neural spiking is modulated by cortical oscillations. To match SWS durations, we used the first 10 min of SWS_{pre} and SWS_{post} for this analysis (Online Methods). Notably, we found that TR_d and TR_i units consistently experienced a significant increase in the peak-to-peak amplitude of the STA in comparison to TU units even in the absence of any changes in the LFP power (**Fig. 2**). This appeared to be the result of a greater likelihood of spikes occurring at a specific phase of the SWA (**Fig. 2a**). Although TR_d and TR_i units experienced 48.70 ± 5.78% and 53.53 ± 3.78% increases in STA amplitude, respectively, TU units experienced a significant net −11.44 ± 3.44% reduction (one-way ANOVA, $F_{2,171} = 52$,

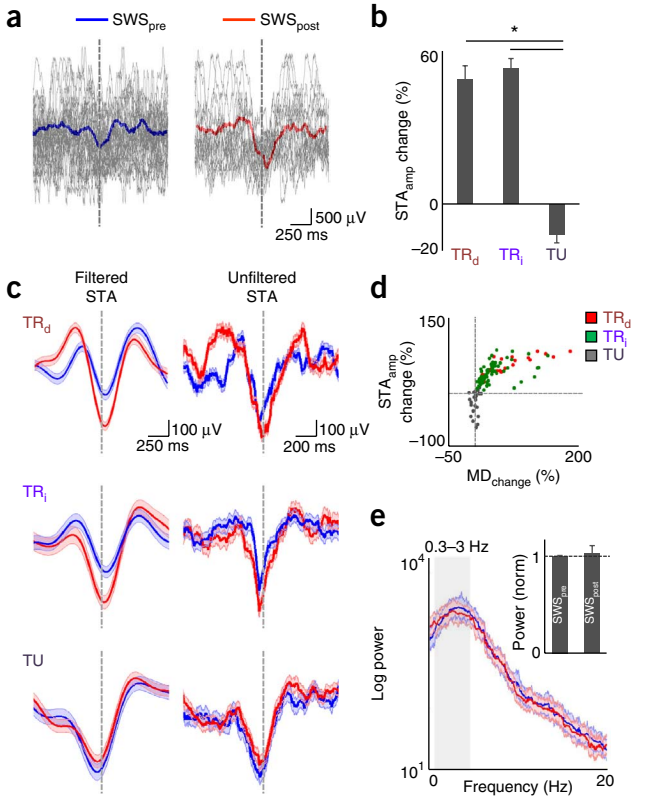

Figure 2 Changes in phase-locking and coherent spiking during SWS after learning. (**a**) Examples of 50 raw LFP traces during SWS_{pre} and SWS_{post} for a task-related direct neuron. Superimposed is the mean trace from an SWS epoch. Each trace was aligned to time of spike (dashed line). (**b**) Percent change in STA amplitude for each of the three categories of units in 15 sessions (mean ± s.e.m., one-way ANOVA, $F_{2,171} = 52$, $P < 10^{-17}$; *post hoc t* test showed significant difference for the TR_d and TU, and TR_d and TU groups, $^*P < 0.05$). (**c**) Filtered (0.3–3 Hz) and unfiltered STAs during SWS before and after successful acquisition of neuroprosthetic skill (shaded areas are s.e.m.). (**d**) Scatter plot showing relationship of depth of modulation during task performance to change in STA amplitude ($R = 0.85$, Spearman correlation, $P < 0.05$). Color code represents each class of neuron. Dashed lines are the 0 values for the *x* and *y* axes. (**e**) Comparison of the power spectra from 10 min of SWS_{pre} and SWS_{post} from a single experiment (shaded area is s.e.m.). Inset compares the power in the 0.3–3-Hz band for multiple experiments ($n = 15$, normalized to SWS_{pre} for each experiment; average change of 3.11 ± 6.48% (mean ± s.e.m.), *t* test, $t_{14} = -0.83$, $P = 0.79$).

$P < 10^{-17}$, *post hoc t* test showed significant difference for TR_d and TU; and TR_d and TU groups, $P < 0.05$; **Fig. 2b**). In contrast, we did not find any significant differences when we calculated the STA for spindle (8–20 Hz) and ripple (100–300 Hz) frequency bands (one-way ANOVA, $F_{2,171} = 0.35$, $P = 0.69$ for spindle-band comparisons; one–way ANOVA, $F_{2,171} = 0.27$, $P = 0.75$ for ripple band; **Supplementary Fig. 4**). We next assessed whether changes in the power or frequency of SWS or firing properties of units could account for our observations. There were no significant changes in either the SWA power (average change of 3.11 ± 6.48%, *t* test, $t_{14} = -0.83$, $P = 0.79$; **Fig. 2e**), density of delta waves in SWS (11.82 ± 5.67 waves min^{-1} versus 12.86 ± 6.09 delta waves min^{-1}, *t* test, $t_{14} = -1.48$, $P = 0.16$; **Supplementary Fig. 5c**), firing rate (6.42 ± 0.43 Hz versus 7.01 ± 0.34 Hz, *t* test, $t_{173} = -1.01$, $P = 0.27$; **Supplementary Fig. 1e–g**) or changes in bursting properties (Kolmogorov-Smirnov test, $P > 0.05$ for over 90% units; **Supplementary Fig. 5a,b**). We also did not find any evidence for a significant relationship between the changes in the firing rates of individual units and changes in STA amplitude (linear regression $R^2 = 0.01$, $P = 0.29$; **Supplementary Fig. 1e**) or the changes in the LFP power relative to change in STA amplitude (linear regression $R^2 = 0.033$, $P = 0.64$). Moreover, we did not find evidence for spatiotemporal changes in the occurrence and rate of SWA (**Supplementary Fig. 6**).

To further establish a link between the task-dependent modulation that emerges after learning and the change in phase-locking to SWA, we examined the relationship between the extent of task-dependent modulation and the percent changes in the STA amplitude. Notably, we found that the degree of TR firing rate modulation significantly predicted the extent of increased STA amplitude during the subsequent SWS (Spearman correlation, $r_{172} = 0.85$, $P < 0.05$; **Fig. 2d**). The non-linear nature of the curve is likely the result of the significant reduction in STA amplitude without a mirror symmetric reduction in depth of modulation for TU units. Consistent with the notion of a link between the extent of task-dependent modulation and the change in STA amplitude, we found that TR_d^+ units experienced a significantly greater change in the STA amplitude than TR_d^- units (*t* test, $t_{25} = 3.37$, $P < 0.005$; **Supplementary Fig. 2d**). Together, these results suggest that there is greater locking of the TR cortical spiking to SWA after successful learning of direct cortical control.

To further confirm that changes in STA amplitude reflect greater phase-locking of TR_d and TR_i units, we used spike-field coherence (SFC) to measure coherence in the 0.3–3-Hz band between the spiking activity and the LFP. The SFC measure complements the STA analysis in that it is not sensitive to changes in power and can readily assess all frequencies. The magnitude of SFC, which varies as a function of frequency and yields a value between 0 and 1, showed a similar trend of enhanced SFC for TR_d and TR_i units in the SWA frequency range (**Fig. 3** and **Supplementary Fig. 7**). Consistent with the STA results in the spindle/ripple bands, there were no detectable differences for other frequency bands (that is, multitaper method using jackknife-based confidence intervals). The TR_d, TR_i and TU groups experienced a change in SFC magnitude of 53.80 ± 7.49%, 60.29 ± 3.61% and −20.21 ± 4.04%, respectively. Thus, TR units (both TR_d and TR_i) showed a significant increase in SFC magnitude compared with TU units (one-way ANOVA, $F_{2,171} = 77$, $P < 10^{-23}$, *post hoc t* test showed significant difference for TR_d and TU; TR_d and TU groups, $P < 0.05$). There were no significant changes in SFC phases of the TR units between SWS_{pre} and SWS_{post} (0.18 ± 0.14 radians, mean ± s.e.m., net change in phase and $P = 0.44$, Watson-Williams circular *t* test

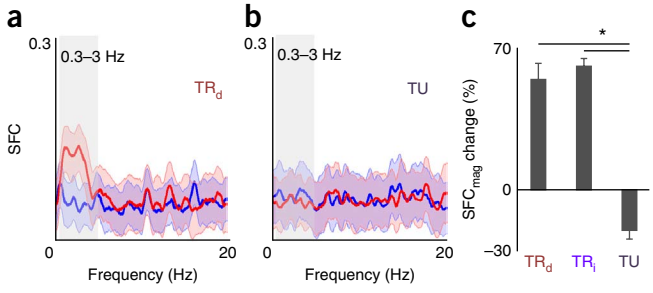

Figure 3 Changes in SFC after learning. (**a**) Example plot of SFC as a function of frequency before (blue) and after (red) skill acquisition for a direct pair. The lighter band is the jackknife error. The box highlights the 0.3–3-Hz band. (**b**) Task-unrelated SFC. (**c**) Mean changes in SFC for the various categories of units (mean ± s.e.m., one-way ANOVA, $F_{2,171} = 77$, $P < 10^{-23}$, *post hoc t* test showed significant difference for TR_d and TU, and TR_d and TU groups, $^*P < 0.05$).

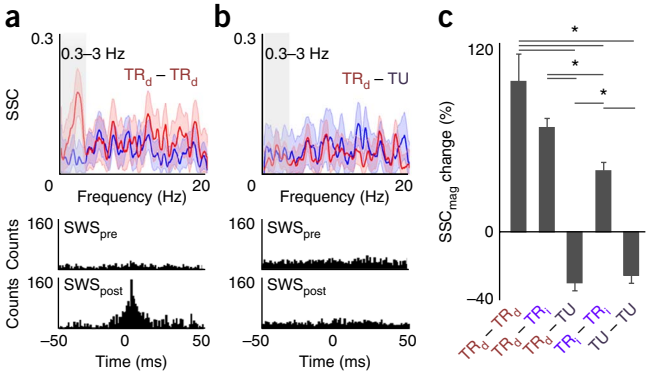

Figure 4 Changes in SSC after learning. (**a**) Example plot of SSC as a function of frequency before (blue) and after (red) skill acquisition for a direct pair. The lighter band is the jackknife error. The box highlights the 0.3–3-Hz band. Below are the respective cross-correlograms from SWS_{pre} and SWS_{post}. (**b**) SSC of a task-unrelated pair. (**c**) Mean changes in SSC for the various categories of neuron pairs are shown as mean ± s.e.m. (one-way ANOVA, $F_{5,271} = 45$, $P < 10^{-29}$; *post hoc t* test, $^*P < 0.05$).

for comparison of phases pre- and post-learning; **Supplementary Fig. 7**). Together with the findings from the STA analysis, these results confirm that TR units, in contrast with TU units, were significantly more phase-locked to SWA after successful learning.

Increased spike-spike coherence after learning

We subsequently assessed whether there are changes in the functional connectivity among the recorded M1 neural ensembles during SWS_{post}. We calculated the magnitude of spike-spike coherence (SSC) for both SWS_{post} and SWS_{pre} for direct units (TR_d) relative to all other units (that is, $TR_d - TR_d$, $TR_d - TR_i$, and $TR_d - TU$ neuronal pairs). We used a 'shuffling' method to assess significant SSC increases in SWS_{post} in comparison to SWS_{pre}. We observed that only the SWS_{post} SSC curves of task related unit pairs showed a significant increase in the 0.3–3-Hz band (**Fig. 4a**,**b**). We found surprisingly robust increases in SSC (0.3–3-Hz range) for TR pairs (that is, $TR_d - TR_d$ and $TR_d - TR_i$) after successful learning (93.15 ± 16.05% increase for 12 $TR_d - TR_d$ pairs and 64.68 ± 4.98% increase for 108 $TR_d - TR_i$ pairs; **Fig. 4c**). In comparison, the change in SSC for the $TR_d - TU$ pairs was −31.41 ± 5.00% (**Fig. 4c**, one–way ANOVA, $F_{2,156} = 65$, $P < 10^{-20}$; *post hoc t* test showed significant differences between $TR_d - TR_d$, and $TR_d - TU$ pairs; and $TR_d - TR_i$, and $TR_d - TU$ pairs, $P < 0.05$). When we used cross–correlation analysis, we also found similar results (one–way ANOVA, $F_{2,156} = 27$, $P < 10^{-10}$; **Fig. 4a**,**b** and **Supplementary Fig. 8**). This indicates that TR units are significantly more likely to fire synchronously than TU units during SWS_{post}.

Lack of changes after poor learning

We next determined whether these changes in STA, SFC and SSC are specific to successful learning. We analyzed four sessions in which animals performed the BMI task, but did not demonstrate robust learning (**Figs. 1e** and **5a**). During these sessions, many of the trials ended because the 15-s timeout was reached. The respective times to completion were 12.22 ± 0.74 s for early trials versus 11.20 ± 0.79 s for late trials (*t* test, $t_3 = 1.7$, $P = 0.18$; **Fig. 1e**), and the percentage of unsuccessful trials remained unchanged (52.50 ± 3.29% early versus 54.17 ± 6.36%, for late; *t* test, $t_3 = -0.52$, $P = 0.68$). In these sessions, TR units did not experience an increase in STA amplitude (**Fig. 5a**,**b**). The mean STA amplitude change for TR_d units was −12.13 ± 3.90% ($n = 8$ units). In addition, the SSC between $TR_d - TR_d$ pairs did not increase in a manner similar to that during robust learning sessions (−15.50 ± 10.32%, $n = 4$ pairs). These were significantly different from the corresponding values in robust learning sessions (*t* test, $t_{14} = 3.9$, $P < 0.01$). Moreover, we also performed control experiments in which rodents were placed in the BMI chamber, but did not perform any task (control sleep; **Fig. 5b**). Instead, they received an equivalent amount of water reward over a time period that matched a typical practice session. The STA before and after these control sessions did not show any increase. The overall change was significantly different from TR_d units from robust learning sessions ($n = 3$ sessions with 32 units, −10.49 ± 2.74% STA change, one–way ANOVA, $F_{2,64} = 59$, $P < 10^{-14}$ with *post hoc t* test showing a significant difference, $P < 0.05$; **Fig. 5b**). This suggests that factors such as sustained attention, reward and attempts at execution without evidence of learning are not sufficient to trigger increased coherent activation during the SWA_{post}. Together, this further supports a conclusion that enhanced locking occurs only for task-related ensembles linked to learning.

Reactivation of emergent task-related ensembles

The analysis presented above examined changes in spike-field relationships and pairwise interactions among TR units after robust and poor learning. We next used a method previously developed to identify task-related neuronal cell assemblies and their coordinated reactivation[35–37]. This method uses principal components analysis (PCA, Online Methods) to identify significant 'signal components', or ensembles of neurons that become temporally correlated during learning and task performance. The output of this analysis are principle components (PCs), consisting of an array of weights assigned to each unit in the identified ensemble, and the eigenvalue, a numerical value that represents the extent of total variance that is captured by a given PC; PCs with the largest eigenvalues capture the most variance (**Fig. 6a**,**b**). Whether a calculated PC represents a significant temporally correlated pattern of activity is determined by λ_{max}, that is, the highest eigenvalue that arises out of an equivalently sized random matrix based on the Marchenko-Pastur law[35–37] (**Fig. 6b**). Thus, PCs with eigenvalues greater than λ_{max} were considered signal components, whereas those below λ_{max} were considered as arising from chance interactions. We started by estimating the signal components (that is, ensemble patterns of activity) linked to successful learning (**Fig. 6a**,**b**). Notably, we found evidence for significant signals only for robust learning sessions (**Fig. 6b**). In contrast, poor-learning sessions generated PCs that were never greater than λ_{max}.

We then examined whether there was a change in the 'reactivation strength' during SWS_{post} by calculating the instantaneous reactivation

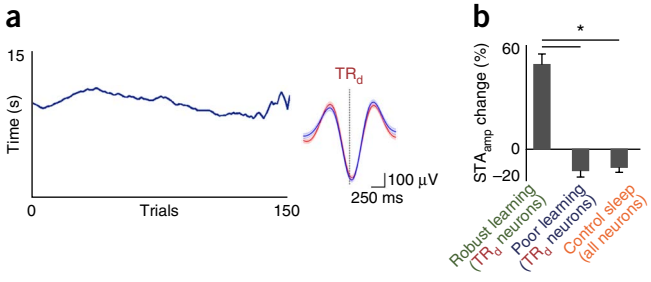

Figure 5 Lack of increase in coherent spiking with poor learning and in control sleep. (**a**) Plot of trial times versus trial number. Right, the STA before and after (shaded area is s.e.). (**b**) Mean changes in direct unit STA for robust learning (reproduced from **Fig. 2b**) and poor learning sessions. Also shown are units from control awake periods (mean ± s.e.m., one-way ANOVA, $F_{2,64} = 59$, $P < 10^{-14}$; significant *post hoc t* tests, $^*P < 0.05$).

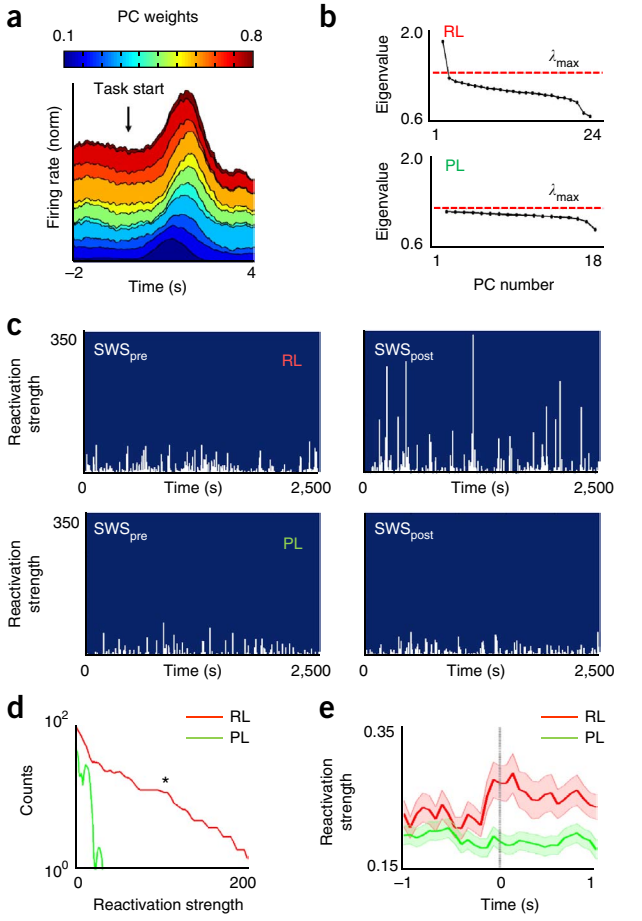

Figure 6 Changes in reactivation strength after learning. (**a**) Peri-event firing rate modulation of units from a single session. Only TR units are shown and sorted by weights in the first PC. (**b**) Correlation matrix eigenvalues calculated from activity during task performance (RL, robust learning; PL, poor learning). Dashed line is the signal threshold (λ_{max}), defined as the theoretical upper bound for a randomized spike train. (**c**) Reactivation strength during SWS_{pre} and SWS_{post} for RL and PL. (**d**) Mean population differences between SWS_{pre} and SWS_{post} distributions of reactivation strengths for RL ($n = 15$ sessions) and PL ($n = 4$ sessions). $^*P < 0.05$ logrank test. (**e**) Event-triggered average of reactivation strength centered on maximum delta wave negativity (time = 0) for SWS_{post} (shaded area represents s.e.m.).

strength of signal components[35–37]. The reactivation strength is an instantaneous measure of how similar the ensemble activity during the SWS epoch is to that identified during the awake period using the PCA analysis. Notably, by comparing the SWS_{pre} and the SWS_{post} epochs, we can identify changes that are specifically linked to learning control. After robust learning sessions, the observed strength of reactivation of the signal component was greatly enhanced in comparison to SWS_{pre} (**Fig. 6c,d**). Moreover, this was not evident during the awake spontaneous periods before the onset of SWS ($n = 15$, logrank test, $P < 0.05$; **Supplementary Fig. 9**). Across multiple sessions, the reactivation strength of the task-related ensemble during SWS_{post} was significantly greater after robust learning than after poor learning (logrank test, $P < 0.05$; **Fig. 6d**). We also examined the specific relationship between instantaneous reactivation and SWA. We therefore performed a delta wave–triggered averaging of reactivation strength. We found that cell assembly reactivation events occurred in concert with maximum delta wave negativity (**Fig. 6e**). The ratio of peak reactivation at the time of maximum delta negativity to the baseline reactivation strength was significantly greater after robust learning than either after poor learning (t test, $t_{17} = -3.60$, $P < 0.005$; **Fig. 6e**) or during spontaneous awake periods after robust learning (t test, $t_{14} = -7.37$, $P < 10^{-5}$; **Supplementary Fig. 9c,d**).

Improvements in task performance

We also conducted experiments in which the animals performed the task during two sessions (referred to as $Block_1$ and $Block_2$) using the same decoder, but separated by a period of spontaneous activity and sleep (**Fig. 7a**). Notably, we found that task performance consistently improved at the start of $Block_2$ ($P < 0.05$ for each of the eight individual comparisons of the last 30 trials from $Block_1$ and the first 30 trials from $Block_2$, overall t test, $t_7 = 4.87$, $P < 0.005$; **Fig. 7b**). Such improvements were also evident if we compared the best performance at the end of $Block_1$ with that of early $Block_2$ (t test, $t_7 = 4.38$, $P < 0.005$; **Supplementary Fig. 10**). That increased motivation after rest does not account for this finding is supported by the lack of significant

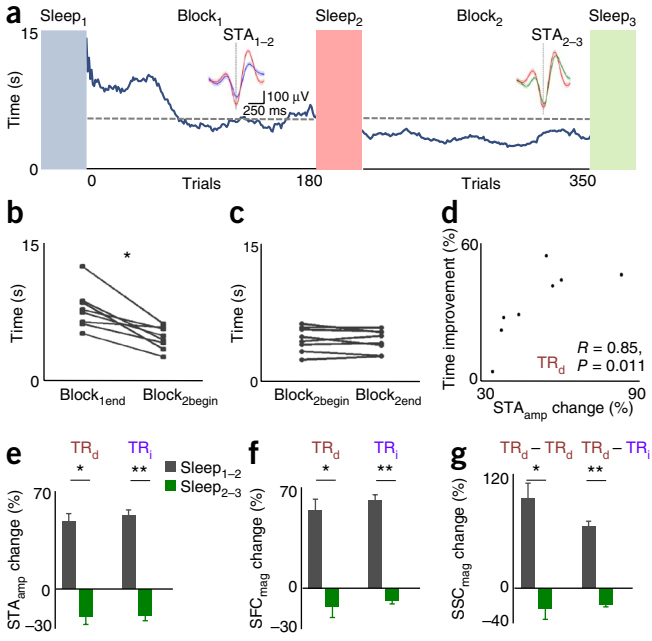

Figure 7 Continued task performance after sleep. (**a**) Plot of time to task completion versus trial number. Data are presented as in **Figure 1c**. Trace inset shows the STA from the respective sleep for a task-related neuron. STA scales are the same (shaded area represents s.e.). (**b**) Average time to reward at the end of first training session ($Block_{1end}$) compared with the beginning of second session ($Block_{2begin}$). $^*P < 0.05$ for all eight comparisons (overall t test, $t_7 = 4.87$, $P < 0.005$). (**c**) Average trial time at the beginning of second session ($Block_{2begin}$) compared with its end ($Block_{2end}$) (overall t test, $t_7 = 0.87$, $P = 0.41$). (**d**) Relationship between mean change in STA amplitude for TR_d units and the improvement in performance. Also noted are the correlation coefficient and P value for the Spearman test. (**e**) Relative changes in STA amplitude for TR_d and TR_i units between $Sleep_{1-2}$ and $Sleep_{2-3}$. STA changes reduced to for both categories significantly (mean ± s.e.m.; TR_d, t test, $t_{40} = 7.77$, $^*P < 10^{-8}$; TR_i, t test, $t_{153} = 11.61$, $^{**}P < 10^{-22}$). (**f**) Relative change in SFC magnitude (mean ± s.e.m.). SFC magnitudes were also reduced significantly (TR_d, t test, $t_{40} = 5.84$, $^*P < 10^{-6}$; $^{**}TR_i$, t test, $t_{153} = 12.20$, $P < 10^{-23}$). (**g**) Relative change in SSC magnitude (mean ± s.e.m.). SSC magnitudes were also reduced for both $TR_d - TR_d$ pairs and for $TR_d - TR_i$ pairs between $Sleep_{1-2}$ and $Sleep_{2-3}$ ($TR_d - TR_d$ pairs, t test, $t_{17} = 5.25$, $^*P < 10^{-4}$; $TR_d - TR_i$ pairs, t test, $t_{153} = 10.45$, $^{**}P < 10^{-18}$).

performance changes during $Block_2$ itself ($P > 0.05$ for each of the eight comparisons, overall t test, $t_7 = 0.87$, $P = 0.41$; **Fig. 7c**). Moreover, although the total period was the same between the two blocks, the time spent in SWS after $Block_1$ was positively correlated with the extent of improvement at the beginning of $Block_2$ (Spearman correlation, $r(6) = 0.67$, $P < 0.05$). Notably, the extent of improvement was also positively correlated with the mean change in STA amplitude between $Sleep_1$ and $Sleep_2$ for TR_d units (Spearman correlation, $r(6) = 0.85$, $P < 0.05$; **Fig. 7d**).

Changes with continued task performance

We first compared the STA amplitude and SFC magnitude changes for TR_d and TR_i units ($n = 8$ experiments, 15 TR_d and 47 TR_i units for $Sleep_{1-2}$ and $Sleep_{2-3}$). With continued execution of the task, there was no further increase in coherent spiking for TR units (**Fig. 7a**,**e**–**g**). STA amplitude changes were $-19.31 \pm 5.94\%$ for TR_d units and -18.63 ± 3.74 for TR_i units (**Fig. 7e**), whereas SFC magnitude changes were $-12.59 \pm 7.7\%$ for TR_d units and -8.71 ± 2.26 for TR_i units (**Fig. 7f**), after the second training session (for $Sleep_{2-3}$). We also examined SSC magnitude changes for TR pairs ($TR_d - TR_d$ and TR_d–TR_i). The changes in SSC magnitude were $-21.37 \pm 11.16\%$ for $TR_d - TR_d$ pairs ($n = 7$) and -16.51 ± 2.89 for $TR_d - TR_i$ pairs ($n = 47$ pairs; **Fig. 7g**). The small changes observed between $Sleep_2$ and $Sleep_3$ (STA_{amp} and SSC_{mag}; **Fig. 7e**–**g**) were not significantly different from the absolute changes seen during poor learning sessions (STA_{amp} changes: TR_d, t test, $t_{21} = -0.85$, $P = 0.40$; TR_i, t test, $t_{53} = -0.70$, $P = 0.48$; SSC_{mag} changes: TR_d –TR_d pairs, t test, $t_9 = -0.38$, $P = 0.71$; $TR_d - TR_i$ pairs, t test, $t_{49} = -0.09$, $P = 0.92$). Together, this indicates that the continued execution of the task in the absence of new learning is not sufficient to trigger further increases in phase-locking and coherent spiking.

DISCUSSION

In summary, we found that successful learning of direct neural control of an actuator is linked to enhanced phase-locking and coherent activation of emergent task-related activity, both direct and indirect, during the post-learning SWS. During the process of learning, these two sets of TR units became increasingly activated with similar temporal profiles. Notably, we found strong evidence for coherent reactivation of these newly functionally coupled ensembles during SWA. In contrast, task-unrelated activity either remained unchanged or experienced a reduction in phase-locking and coherent activation. For control sleep sessions, poor-learning sessions and second sessions without evidence of new performance gains, we did not find evidence of increased phase-locking to SWA or coherent activation, indicating specificity of our observed findings to new learning and skill acquisition. Notably, we found a positive correlation between both time spent in SWS and the extent of STA amplitude change and improvements in task performance following awakening.

Processing of an emergent task-related ensemble

Learning was associated with the emergence of a novel task-related ensemble of both direct and indirect units. In our experiments, randomly chosen pairs of direct units were conditioned through feedback to be volitionally modulated to move an actuator. Consistent with a growing body of literature[3,33,34], modulation of direct units was accompanied by similar task-related firing rate changes in a subset of indirect units (TR_i). Nearly all direct and a subset of indirect units became increasingly modulated in a task-dependent manner with practice (**Fig. 1**). The emergent functional coupling between these units appeared to be a consequence of learning and successful completion of the neuroprosthetic task. Because direct units are embedded in highly connected M1 cortical networks, it is possible that such modulation of adjacent indirect activity is important for precise cortical control.

During the post-learning SWS, there was significantly higher coincident activation of the emergent task-related ensemble. We also found that the extent of firing-rate modulation during task performance predicted the strength of the modifications evident during SWS (**Fig. 2d**). In the absence of learning or new performance gains, this was not present. This strongly suggests that the functional coactivation of units during successful learning is directly linked to the modified functional connectivity and coherent reactivation detected during SWS. We also noted a positive correlation between time spent in SWS and subsequent improvements in task performance. This observed performance gain is consistent with past studies, suggesting that even brief periods of sleep can improve motor performance[20–24]. It remains unclear whether our observed offline processing during SWA is sufficient and necessary to trigger performance gains following awakening; there may be other processes recruited during SWS that could also contribute to task improvements[17,38].

Broader link to sleep and memory

A growing body of literature has linked offline processing during sleep to memory consolidation[17,18]. Memory formation in the hippocampal system is the most widely studied[19,39–42]. In general, after an initial encoding phase, hippocampus-based memories appear to undergo a process of consolidation in which its representation is stabilized to neocortical systems. Although the precise mechanisms underlying such consolidation are incompletely understood, sleep-dependent interactions between hippocampal and cortical circuits have been linked to this process of consolidation. Spontaneous reactivation of neurons reflecting previous experiences have been found during both non-REM (that is, including SWA) and REM sleep in both the hippocampus and cortex[19]. Recent studies have also suggested that both disruption of hippocampal circuits with electrical stimulation and sensory cue–dependent facilitation of memory formations can impede and facilitate memory formation, respectively, suggesting a direct link between the offline processing and the consolidation of episodic memories[43,44].

Sleep, in general, and SWS, in particular, also appear to be important for the consolidation of procedural memories[21–25]. For example, after learning a new motor skill, offline processing during overnight sleep[20–23] or brief daytime naps[24] can facilitate consolidation. Local changes in SWA are also associated with motor learning and performance[26–29]. After new motor learning or sensory stimulation there appear to be local increases of SWA in specific cortical areas or hemispheres[26,45]. In contrast, disuse and inactivity can locally decrease SWA and deteriorate task performance[28]. It remains incompletely understood, however, how neural ensembles in motor cortex are precisely modified during SWA after skill acquisition. In this context, BMIs offer a powerful tool to directly modulate neurons in motor cortex, thereby allowing precise characterization of how both task-related and task-unrelated ensembles are differentially processed in motor cortex during SWA. Although the exact link between neuroprosthetic learning and general procedural learning remains unclear, our results suggest that the emergent task-related ensembles could be also processed in a similar manner after natural motor learning.

Implications for neuroprosthetic control

An important goal of the field of BMIs is to allow stable control of complex devices over long periods of time. Although decoder adaptation can speed the overall rate of skill acquisition, adaptation of neural circuits and neural plasticity over longer periods of time

may be essential for stable skill acquisition. This may be especially important for skilled control over complex devices that resemble our natural control of limbs. For example, a recent study highlighted that optimal recruitment of long-term neural plasticity is essential for achieving flexible control that resembles our natural abilities[14]. The study illustrated that long-term neural plasticity (that is, across multiple sessions) is essential for achieving BMI control that can readily switch between two control schemes without interference. Specifically, only after long-term stabilization of a single control scheme was simultaneous acquisition of a second control scheme possible. Based on the parallels between motor and neuroprosthetic control, a reasonable hypothesis is that consolidation of 'prosthetic memory' is essential for such apparent resistance to interference. Our observed phenomenon of offline processing during SWS may consolidate such a prosthetic memory.

Conclusion

Our results provide direct evidence that an emergent group of TR units, which were increasingly functionally coupled with learning and task performance, experienced coherent reactivation during SWS_{post}. Knowledge about this phenomenon should help to better account for mechanisms of neural plasticity and further the goal of achieving stable long-lasting neuroprosthetic control that resembles natural motor control.

METHODS

Methods and any associated references are available in the online version of the paper.

Note: Any Supplementary Information and Source Data files are available in the online version of the paper.

ACKNOWLEDGMENTS

This work was supported by the Department of Veterans Affairs (B6674 to K.G.), the Burroughs Wellcome Fund (1009855 to K.G.), the American Heart/Stroke Association (0875016N to K.G.) and a VA Psychiatric Research Advanced Fellowship (to D.S.R.).

AUTHOR CONTRIBUTIONS

C.C.W., T.G. and K.G. conducted the experiments. T.G. and D.S.R. analyzed the data. K.G. supervised the study. All of the authors contributed to the writing and editing of the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

1. Chapin, J.K., Moxon, K.A., Markowitz, R.S. & Nicolelis, M.A. Real-time control of a robot arm using simultaneously recorded neurons in the motor cortex. *Nat. Neurosci.* **2**, 664–670 (1999).
2. Koralek, A.C., Jin, X., Long, J.D. II, Costa, R.M. & Carmena, J.M. Corticostriatal plasticity is necessary for learning intentional neuroprosthetic skills. *Nature* **483**, 331–335 (2012).
3. Arduin, P.J., Fregnac, Y., Shulz, D.E. & Ego–Stengel, V. "Master" neurons induced by operant conditioning in rat motor cortex during a brain-machine interface task. *J. Neurosci.* **33**, 8308–8320 (2013).
4. Carmena, J.M. *et al.* Learning to control a brain-machine interface for reaching and grasping by primates. *PLoS Biol.* **1**, E42 (2003).
5. Serruya, M.D., Hatsopoulos, N.G., Paninski, L., Fellows, M.R. & Donoghue, J.P. Instant neural control of a movement signal. *Nature* **416**, 141–142 (2002).
6. Taylor, D.M., Tillery, S.I. & Schwartz, A.B. Direct cortical control of 3D neuroprosthetic devices. *Science* **296**, 1829–1832 (2002).
7. Santhanam, G., Ryu, S.I., Yu, B.M., Afshar, A. & Shenoy, K.V. A high-performance brain-computer interface. *Nature* **442**, 195–198 (2006).
8. Moritz, C.T., Perlmutter, S.I. & Fetz, E.E. Direct control of paralyzed muscles by cortical neurons. *Nature* **456**, 639–642 (2008).
9. Jarosiewicz, B. *et al.* Functional network reorganization during learning in a brain-computer interface paradigm. *Proc. Natl. Acad. Sci. USA* **105**, 19486–19491 (2008).
10. Musallam, S., Corneil, B.D., Greger, B., Scherberger, H. & Andersen, R.A. Cognitive control signals for neural prosthetics. *Science* **305**, 258–262 (2004).
11. Hochberg, L.R. *et al.* Neuronal ensemble control of prosthetic devices by a human with tetraplegia. *Nature* **442**, 164–171 (2006).
12. Hochberg, L.R. *et al.* Reach and grasp by people with tetraplegia using a neurally controlled robotic arm. *Nature* **485**, 372–375 (2012).
13. Collinger, J.L. *et al.* High-performance neuroprosthetic control by an individual with tetraplegia. *Lancet* **381**, 557–564 (2013).
14. Ganguly, K. & Carmena, J.M. Emergence of a stable cortical map for neuroprosthetic control. *PLoS Biol.* **7**, e1000153 (2009).
15. Gilja, V. *et al.* A high-performance neural prosthesis enabled by control algorithm design. *Nat. Neurosci.* **15**, 1752–1757 (2012).
16. Stickgold, R. Sleep-dependent memory consolidation. *Nature* **437**, 1272–1278 (2005).
17. Diekelmann, S. & Born, J. The memory function of sleep. *Nat. Rev. Neurosci.* **11**, 114–126 (2010).
18. Peigneux, P., Laureys, S., Delbeuck, X. & Maquet, P. Sleeping brain, learning brain. The role of sleep for memory systems. *Neuroreport* **12**, A111–A124 (2001).
19. Lee, A.K. & Wilson, M.A. Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* **36**, 1183–1194 (2002).
20. Walker, M.P., Brakefield, T., Morgan, A., Hobson, J.A. & Stickgold, R. Practice with sleep makes perfect: sleep-dependent motor skill learning. *Neuron* **35**, 205–211 (2002).
21. Fischer, S., Nitschke, M.F., Melchert, U.H., Erdmann, C. & Born, J. Motor memory consolidation in sleep shapes more effective neuronal representations. *J. Neurosci.* **25**, 11248–11255 (2005).
22. Cohen, D.A., Pascual–Leone, A., Press, D.Z. & Robertson, E.M. Off-line learning of motor skill memory: a double dissociation of goal and movement. *Proc. Natl. Acad. Sci. USA* **102**, 18237–18241 (2005).
23. Fischer, S., Hallschmid, M., Elsner, A.L. & Born, J. Sleep forms memory for finger skills. *Proc. Natl. Acad. Sci. USA* **99**, 11987–11991 (2002).
24. Korman, M. *et al.* Daytime sleep condenses the time course of motor memory consolidation. *Nat. Neurosci.* **10**, 1206–1213 (2007).
25. Walker, M.P., Brakefield, T., Hobson, J.A. & Stickgold, R. Dissociable stages of human memory consolidation and reconsolidation. *Nature* **425**, 616–620 (2003).
26. Huber, R., Ghilardi, M.F., Massimini, M. & Tononi, G. Local sleep and learning. *Nature* **430**, 78–81 (2004).
27. Tamaki, M. *et al.* Enhanced spontaneous oscillations in the supplementary motor area are associated with sleep-dependent offline learning of finger-tapping motor-sequence task. *J. Neurosci.* **33**, 13894–13902 (2013).
28. Huber, R. *et al.* Arm immobilization causes cortical plastic changes and locally decreases sleep slow wave activity. *Nat. Neurosci.* **9**, 1169–1176 (2006).
29. Hanlon, E.C., Faraguna, U., Vyazovskiy, V.V., Tononi, G. & Cirelli, C. Effects of skilled training on sleep slow wave activity and cortical gene expression in the rat. *Sleep* **32**, 719–729 (2009).
30. Moroni, F. *et al.* Procedural learning and sleep hippocampal low frequencies in humans. *Neuroimage* **42**, 911–918 (2008).
31. Inostroza, M. & Born, J. Sleep for preserving and transforming episodic memory. *Annu. Rev. Neurosci.* **36**, 79–102 (2013).
32. Ganguly, K., Dimitrov, D.F., Wallis, J.D. & Carmena, J.M. Reversible large-scale modification of cortical networks during neuroprosthetic control. *Nat. Neurosci.* **14**, 662–667 (2011).
33. Fetz, E.E. Volitional control of neural activity: implications for brain-computer interfaces. *J. Physiol.* (*Lond.*) **579**, 571–579 (2007).
34. Koralek, A.C., Costa, R.M. & Carmena, J.M. Temporally precise cell-specific coherence develops in corticostriatal networks during learning. *Neuron* **79**, 865–872 (2013).
35. Peyrache, A., Khamassi, M., Benchenane, K., Wiener, S.I. & Battaglia, F.P. Replay of rule-learning related neural patterns in the prefrontal cortex during sleep. *Nat. Neurosci.* **12**, 919–926 (2009).
36. Peyrache, A., Benchenane, K., Khamassi, M., Wiener, S.I. & Battaglia, F.P. Principal component analysis of ensemble recordings reveals cell assemblies at high temporal resolution. *J. Comput. Neurosci.* **29**, 309–325 (2010).
37. Lopes-dos-Santos, V., Ribeiro, S. & Tort, A.B. Detecting cell assemblies in large neuronal populations. *J. Neurosci. Methods* **220**, 149–166 (2013).
38. Tononi, G. & Cirelli, C. Sleep and the price of plasticity: from synaptic and cellular homeostasis to memory consolidation and integration. *Neuron* **81**, 12–34 (2014).
39. Marshall, L. & Born, J. The contribution of sleep to hippocampus-dependent memory consolidation. *Trends Cogn. Sci.* **11**, 442–450 (2007).
40. Pavlides, C. & Winson, J. Influences of hippocampal place cell firing in the awake state on the activity of these cells during subsequent sleep episodes. *J. Neurosci.* **9**, 2907–2918 (1989).
41. Wilson, M.A. & McNaughton, B.L. Reactivation of hippocampal ensemble memories during sleep. *Science* **265**, 676–679 (1994).
42. Buzsaki, G. Two-stage model of memory trace formation: a role for 'noisy' brain states. *Neuroscience* **31**, 551–570 (1989).
43. Rasch, B., Buchel, C., Gais, S. & Born, J. Odor cues during slow-wave sleep prompt declarative memory consolidation. *Science* **315**, 1426–1429 (2007).
44. Jadhav, S.P., Kemere, C., German, P.W. & Frank, L.M. Awake hippocampal sharp-wave ripples support spatial memory. *Science* **336**, 1454–1458 (2012).
45. Vyazovskiy, V., Borbely, A.A. & Tobler, I. Unilateral vibrissae stimulation during waking induces interhemispheric EEG asymmetry during subsequent sleep in the rat. *J. Sleep Res.* **9**, 367–371 (2000).

ONLINE METHODS

Animals and surgery. Experiments were approved by the Institutional Animal Care and Use Committee at the San Francisco VA Medical Center. We used six adult Long-Evans male rats that were approximately 3 months old. No statistical test was run to determine sample size a priori. The sample sizes chosen are similar to those used in previous publications. Animals were kept under controlled temperature and a 12-h light:12-h dark cycle with lights on at 06:00 a.m. Probes were implanted during a recovery surgery performed under isofluorane (1–3%) anesthesia following induction with ketamine (2 mg per kg of body weight) and xylazine (1 mg per kg). Atropine sulfate was also administered before anesthesia (0.02 mg per kg). The post-operative recovery regimen included administration of buprenorphine at 0.02 mg per kg and meloxicam at 0.2 mg per kg. Dexamethasone at 0.5 mg per kg and trimethoprim sulfadiazine at 15 mg per kg were also administered post-operatively for 5 d. We used 32-channel microwire arrays (33-μm polyimide-coated tungsten microwire arrays) in four rats. We also used 32-channel silicon probe arrays in a tetrode configuration in two rats. Arrays were lowered down to 1,400–1,800 μm in the primary motor cortex (M1) in the upper limb area (1–3 mm anterior to bregma and 2–4 mm lateral from midline). The reference wire was wrapped around a screw inserted in the midline over the cerebellum. Final localization of depth was based on quality of recordings across the array at the time of implantation. In one rat, we implanted electromyogram (EMG) wires into neck muscles. Specifically, pairs of Teflon-coated stainless steel wires (AM Systems) were inserted in neck muscles and tunneled subcutaneously to a connector on the cap. All animals were allowed to recover for 1 week before start of experiments.

Electrophysiology. We recorded extracellular neural activity using tungsten microwire electrode arrays (MEAs, $n = 4$ rats, Tucker-Davis Technologies) or silicon probes in a tetrode configuration ($n = 2$ rats, Neuronexus Technologies). We recorded spike and LFP activity using a 128-channel TDT-RZ2 system (Tucker-Davies Technologies). Spike data was sampled at 24,414 Hz and LFP data at 1,018 Hz. ZIF-clip based analog headstages with a unity gain and high impedance (~1 GΩ) were used. Differential EMG was also recorded at 1,018 Hz and high-pass filtered with low-frequency cut off at 100 Hz. Only clearly identifiable units with good waveforms and high signal to noise were used. The remaining neural data was recorded for offline analysis. Behavior related timestamps (trial onset, trial completion) were sent to the RZ2 analog input channel using an Arduino digital board and synchronized to neural data.

We used the term 'unit' to refer to the sorted spike recordings from both the MEA and tetrode recordings. For both, we initially used an online sorting program (SpikePac, Tucker-Davies Technologies) for neuroprosthetic control (either in single channel or tetrode mode). We then conducted offline sorting. We sorted the MEA recordings using standard offline cluster cutting methods in Tucker-Davies Technologies' OpenSorter software. We subsequently used waveform shape and the presence of an absolute/relative refractory period in the interspike interval to judge quality of isolation (**Supplementary Fig. 1**)[14]. Offline tetrode sorting was performed using a previously described method (**Supplementary Fig. 1**)[46,47]. Specifically, a voltage-based threshold was set based on visual inspection for each channel that allowed for best separation between putative spikes and noise; typically this threshold was at least 4 s.d. away from the mean. Events were time-stamped and waveforms for each event were peak aligned. K-means clustering was then performed across the entire data matrix of waveforms (30 samples per channel × 4 channels × number of waveforms). Automated sorting was performed by: (1) first over-clustering waveforms using a K-means algorithm (split into many mini–clusters), (2) calculation of interface energy (a nonlinear similarity metric that allows for an automated decision of whether mini-clusters are actually part of the same cluster), and (3) followed by aggregation of similar clusters. Such aggregation allows for a reduction in the total numbers of clusters that need to be manually inspected. Automated sorting was followed by manual inspection and sorting of spikes (including both merging and splitting clusters further, and removing significant outliers based on Gaussian distribution), using feature space, auto-correlations, cross-correlations and linear discriminant analysis to determine which clusters represent single units and to prevent over-sorting. **Supplementary Figure 1** illustrates examples of sorting achieved with the MEA and silicon probe tetrode recordings.

Behavior. After recovery, animals were typically gentled for several days before the start of experimental sessions. Animals acclimated to a custom plexiglass behavioral box (**Fig. 1a**) during this period. The box was equipped with a slit (covered with a door at one end) that served as a drinking zone. Initially, water delivery from the actuator was not introduced to the rats and they were just acclimatized to the box. Toward the end of the acclimation period, the rats typically fell asleep while in the box. Animals were then water scheduled such that water (from the feeding tube illustrated in **Fig. 1a**) was available in a randomized fashion while in the behavioral box. We monitored body weights on a daily basis to ensure that the weight did not drop below 95% of the initial weight.

Behavioral sessions were typically conducted in the morning, with second sessions conducted in the afternoon. Two single learning sessions were conducted in the afternoon. We found the same results for these experiments. We recorded neural data from the rats for 2 h before start of BMI training. The rats were then allowed to perform the task over a 2-h session. Recorded neural data was entered in real-time to custom routines in Matlab. These then served as control signals for the angular velocity of the feeding tube. The rats typically performed ~180–200 trials for robust learning sessions (for example, **Figs. 1c and 7a**). Following this, we recorded neural data from animals for a 2-h period. For the two session experiments, the animals then continued with another 2-h training session followed by another spontaneous activity recording. Sorted units at the beginning of the recording were checked for maintenance throughout the second training session. We also collected 'control sleep' data in a subset of animals. After an initial spontaneous activity period with sleep, we delivered random automatic water rewards over a period to match a typical time period for task performance. We then recorded spontaneous activity and sleep.

Neural control of the feeding tube. During the BMI training sessions, we typically selected two well-isolated units as 'direct' and allowed their neural activity to control the angular velocity of the feeding tube (**Fig. 1d**). In 3 of the 15 sessions, there was only one neuron selected as the direct unit. These units maintained their stability throughout the recording as evidenced by stability of waveform shape and interspike-interval histograms[14]. We binned the spiking activity into 100-ms bins. We then established a mean firing rate for each neuron over a 3–5-min baseline period. The mean firing rate was then subtracted from its current firing rate at all times. The specific transform that we used was

$$\Theta_{\mathbf{v}} = C * [G_1 * r_1(i) + G_2 * r_2(i)]$$

where $\Theta_{\mathbf{v}}$ is the angular velocity of the feeding tube, $r_1(i)$ and $r_2(i)$ are firing rates of the direct units. G_1 and G_2 are randomized coefficients that ranged from +1 to −1 and were held constant after initialization. C is a fixed constant that scales the firing rates to angular velocity. The animals were then allowed to control the feeding tube via modulation of neural activity. The tube started at the same position at the start of each trial (P_1 in **Fig. 1a**). The calculated angular velocity was added to the previous angular position at each time step (100 ms). During each trial, the angular position could range from −45 to +180 degrees. If the tube stayed in the 'target zone' (P_2 in **Fig. 1a**; spanned 10° area) for a period of 300 ms a water reward was delivered. In the beginning of a session, most rats were unsuccessful at bringing the feeding tube to position P_2. Most rats steadily improved control and reduced the time to completion of the task during the first session. In some cases, animals did not improve control (that is, poor learning shown in **Fig. 5**). In subsequent sessions, these animals did demonstrate that they could learn the task. As shown in **Supplementary Table 1**, multiple learning sessions were obtained from each animal. These sessions were typically 1 week apart to ensure that new units were recorded. Consistent with past studies, we also found that incorporation of new units into the control scheme required new learning[8,14,15]. In many sessions we videotaped the rat during the BMI training blocks. Consistent with multiple reports, we did not observe movements that systematically predicted feeding tube movements[2,15,32]. Specifically, we analyzed whether limb movements measured using the video recording (that is, markers manually assigned to the head, torso and each limb using image processing software) covaried with movements of the feeding tube. Across multiple sessions, we did not find evidence for significant covariation (data not shown). This is likely a result of the fact that nonmovement–related random weights were assigned to the direct units.

doi:10.1038/nn.3759

Sessions and changes in performance. Analysis was performed in Matlab (Mathworks) with custom-written routines. A total of 19 training sessions recorded from 6 rats were used for our initial analysis. 15 of these sessions had robust learning (that is, >3 s.d. drop in time to completion in the last one-third of trials, or late trials, in comparison with the first one-third of trials, or early trials). 8 of these 15 sessions also had a second follow-up training session (**Fig. 7**). Four additional sessions were classified as poor learning (no significant improvement in time to task completion; **Supplementary Table 1**). In addition we analyzed three separate control sleep sessions. As outlined above, during these experiments we matched the time awake using a randomized reward schedule that did not require learning. For **Figure 7b**,**c**, we compared changes in task performance between and across sessions. Specifically, we compared the performance change between $Block_1$ and $Block_2$ by calculating the mean and s.d. of the time to completion during the last 30 trials in $Block_1$ and the first 30 trials $Block_2$ (**Fig. 7b**). We also compared the first and last 30 trials in $Block_2$ (**Fig. 7c**). Furthermore, we also compared the best performance in $Block_1$ (30 trials with the shortest time to reward) and the first 30 trials $Block_2$ (**Supplementary Fig. 10**). We used a paired *t* test to assess statistical significance.

Task-related activity. The distinction between TR_d, TR_i and TU units was based on whether significant modulation of baseline firing activity after the Go cue was present (peak of modulation at the time >2.5 s.d. above the baseline period). For the correlation between task-related modulation and changes in amplitude of the STA, the percentage change in modulation was calculated for firing rates during the first four seconds after the Go cue relative to 3 s of the baseline segment (before the Go cue).

Putative cell-type identification. This was performed by comparing the widths between the first trough in the negatively detected spikes and the first peak after that on the mean waveforms (**Supplementary Fig. 1b–d**).

Identification of NREM oscillations. Identification of pre- and post-SWS epochs was performed by combined visual assessment of low-frequency, high-amplitude slow-wave oscillations as well as a 3 s.d. threshold of the filtered data (0.3–3 Hz)[48,49]. If there was a sustained reduction >1.5 s in the amplitude of the SWA below threshold during a continuous epoch we excluded these segments. For the analysis of STA's and coherence we matched the SWS time periods used for each (first 600 s). For subsequent analysis (that is, Spearman correlation below), we used the entire time spent in SWS per session. Delta waves were detected directly from a filtered, *z* scored SWS in the median LFP channel in each animal (median value of *z* scored LFP from all channels). After identification, only waves that contained a negativity larger than 80 μV (in the raw LFP) were included[49]. These criteria resulted in reproducible and specific detection of high-amplitude delta waves in pre and post SWS epochs. We classified SWS using video analysis in five rats and EMG recordings/video analysis in one rat. We used the EMG recordings to confirm SWS (**Supplementary Fig. 3**). The video recordings were processed using MATLAB (Image Processing Toolbox) to detect any movements. We specifically subtracted a baseline image of the cage environment from each frame, converted the difference into a black and white image, filtered the resulting image for large objects, and then used segmentation analysis to identify the rodent's body/head. We found that the previously described '40-s rule'[50] was a reliable indicator of sleep in comparison to EMG recordings. There were no significant differences between the duration of SWS (*t*–test, $t_{14} = -0.02$, $P = 0.98$) in the SWS_{pre} and SWS_{post} blocks (**Supplementary Fig. 3b**). Moreover, we also examined if there were changes in the frequency and temporal lags between channels with a TR_d and those with 'distant' TU's (greater than two electrodes apart). We did not find any differences in the temporal latencies (*t* test, $t_{13} = -0.30$, $P = 0.76$; **Supplementary Fig. 6b**) and the OFF period density (*t* test, $t_{13} = -0.31$, $P = 0.77$; **Supplementary Fig. 6c**).

STA. We first calculated the STA to measure how spikes locked to the SWA. For the SWS, we used both the filtered (0.3–3 Hz) and the unfiltered LFP. Given the predominance of high-amplitude low frequencies, these were not significantly different (**Fig. 2c**). For the filtering we used a zero-phase distortion filter in MATLAB[49]. Although we used 600 s for comparing the changes shown in the figure, we found that effects were similar even if all recorded SWS durations were used (data not shown). However, for subsequent analysis in the spindle (8–20 Hz) and ripple (100–300-Hz range) bands, we used filtered data to minimize contamination from other bands[49]. During the task, we found that the task-related STA changed only in the beta range (12–40-Hz band, but not in other bands, one-way ANOVA, $F_{2,171} = 36$, $P < 10^{-13}$; **Supplementary Fig. 2e**,**f**). We used the first 5 s after the start of the trial to calculate these STAs. We also used the LFP from a neighboring channel at a distance ranging from 150–250 μm (depending on if it was a MEA or silicon probe). For a subset of the data, we also ensured that coherent spiking did not contribute to our results. Thus, we excised LFP data around the spike times (in the neighboring channel) and linearly interpolated the excised segment using neighboring data points. Our results were identical both with and without this procedure.

Coherency measures. We used the Chronux toolbox to calculate both the SSC and SFC[51] (http://chronux.org/). The magnitude of both is a function of frequency and takes values between 0 and 1. We segmented the pre- and post-SWS into 20-s segments and then averaged the coherency measures across segments. For the multitaper analysis, we used a time-bandwidth (TW) product of 10 with 19 tapers. To compare coherences across groups, a *z* score was calculated using the programs available in the Chronux Toolkit. Coherence between activity in two regions was calculated and defined as

$$Cxy = \frac{\left|R_{xy}\right|}{\sqrt{\left|R_{xx}\right|}\sqrt{\left|R_{yy}\right|}}$$

where R_{xx} and R_{yy} are the power spectra and R_{xy} is the cross-spectrum. Spectral analysis were calculated in segmented SWS epochs and averaged across these epochs across animals. Mean coherence in the slow-wave band was calculated for the 0.3–3-Hz band. Significance testing on coherence estimates was performed on mean slow-wave band estimates across spike-field and spike-spike pairs using one-way ANOVA and independent samples *t* tests for multiple comparisons. Circular statistics were used for comparing phases of the different cell types before and after learning. Specifically, Raleigh test for homogeneity was used to check the presence of a uniform distribution of the phases and Watson-Williams circular *t* test was used to compare changes in phase. The change of magnitude for SFC and SSC were analyzed between the 0.3- and 3-Hz frequency bands. For SSC analysis, task-related direct unit with greatest depth modulation was used to calculate SSC for every other unit. For TU – TU and TR_i – TR_i pairs, the least modulated unrelated unit and greatest modulated indirect unit were chosen, respectively. The power spectrum of the LFP channels used in the coherence calculation, as well as for overall SWA power change in pre- and post-SWS, was also determined using the multitaper method. The power spectrum for the full sleep was not significantly different (*t* test, $t_{14} = -1.95$, $P = 0.07$).

Spike cross-correlation. We computed the cross-correlation histogram (CCH) for the same TR_d – TR_d, TR_d – TR_i and TU – TU pairs that were evaluated during the SSC analyses. The spike counts were equalized in pre- and post-SWS and the pairwise coincident firing was quantified by taking the ratio of the peak (within a time period of ±5 ms) in the center with tail on side (time period of ±10–20 ms; **Supplementary Fig. 8**). For a subset of data we also constructed pseudo random spike train CCH with identical spike counts, mean firing rate and duration to the experimental data. Such simulations were run 200 times (Monte Carlo Simulations). Significant CCH in these cases were the ones with higher central peaks than 99% of the simulated data. Only task related neuronal pairs were found to have significant CCHs.

Ensemble reactivation analyses. Reactivation analysis characterized the neural activity patterns in SWS_{pre} and SWS_{post} using a template that was created during task execution in the awake period[35–37]. We computed a pairwise unit activity correlation matrix as described previously[36] for SWS_{pre}, awake and SWS_{post}. The eigenvector with the largest eigenvalue in the awake correlation matrix served as the learning–related cell assembly. For all epochs, spike trains were binned ($\Delta t_{bin} = 50$ ms). The obtained spike counts for each cell ($s^i(t)$, $i = 1{:}n$, $t = 1{:}B$, where n is the number of cells and B is the number of time bins in the epoch) were *z* transformed, obtaining the Q matrix

$$Q_{it} = \frac{s^i(t) - \langle s^i \rangle}{\sigma_{si}}$$

where $\langle s^i \rangle$ is the mean and σ_{si} is the s.d.

The pairwise cell activity correlation matrix can be presented as

$$C = \frac{1}{B} QQ^t \quad (1)$$

Thus, we obtained three spike count matrices, QSWSpre, Qawake and QSWSpost, and the three correlation matrices, CSWSpre, Cawake and CSWSpost. For quantification of the instantaneous match of activity during SWS to learning related activity during awake periods, we established a C^{template} (awake) and C^{match} (SWS_{pre} and SWS_{post}). As outlined previously[36], we also calculated a measure of similarity, M, using the elements of C^{template} and C^{match}

$$M^{\text{template-match}} = \sum_{i,j,i<j} C_{ij}^{\text{match}} C_{ij}^{\text{template}}$$

$$= \frac{1}{2B^{\text{match}}} \sum_{t=1}^{B^{\text{match}}} \sum_{i,j,i\neq j} Q_{ij}^{\text{match}} C_{ij}^{\text{template}} Q_{jt}^{\text{match}}$$

To distinguish among groups, we calculated the eigenvectors of C^{template}, namely $p^{(l)}$ (with $p^{(1)}$ associated with the largest eigenvalue λ_1). We then decomposed C^{template} into projectors, $P^{(l)}$, defined by the outer products of the eigenvectors with themselves.

$$C^{\text{template}} = \sum_l \lambda_l P^{(l)} = \sum_l \lambda_l p^{(l)} (p^{(l)})^T$$

We also generated a time series $R_l^{\text{match}}(t)$ that measured the instantaneous match of the l-th template, corresponding to the eigenvectors described below. Put another way, this value is the 'reactivation strength' or the instantaneous measure of how similar the ensemble activity during SWS_{pre} and SWS_{post} is to that identified during the awake period.

$$M^{\text{template-match}} = \frac{1}{2B^{\text{match}}} \sum_l \lambda_l \sum_{t=1}^{B^{\text{match}}} \sum_{i,j,i\neq j} Q_{ij}^{\text{match}} P^{(l)} Q_{jt}^{\text{match}}$$

$$= \frac{1}{2B^{\text{match}}} \sum_l \lambda_l \sum_{t=1}^{B^{\text{match}}} R_l^{\text{match}}(t)$$

with

$$R_l^{\text{match}}(t) = \sum_{i,j,i\neq j} Q_{ij}^{\text{match}} P^{(l)} Q_{jt}^{\text{match}}$$

Q_{ij}^{match} is a vector of the multi-cell counts at time point t. The previous study[36] further proposed the use of the Marcenko-Pastur distribution as the null hypothesis for the existence of cell assemblies. It was demonstrated that eigenvalues of the correlation matrix of a normal random matrix R with statistically independent rows follows a probability distribution described.

$$p(\lambda) = \frac{q}{2\pi\sigma^2} \frac{\sqrt{(\lambda_{\max} - \lambda)(\lambda - \lambda_{\min})}}{\lambda}$$

where

$$\lambda_{\min}^{\max} = \sigma^2 (1 \pm \sqrt{1/q})^2$$

And σ^2 is the variance of the elements of the random matrix R, which is 1 here (due to z scoring), and $q = R_{\text{columns}}/R_{\text{rows}} \geq 1$. Under the null hypothesis of an uncorrelated matrix, the correlations between spike trains are determined only by random fluctuations and the eigenvalues of template awake matrix must lie between $\lambda_{\min}$ and $\lambda_{\max}$. Eigenvalues greater than $\lambda_{\max}$ are therefore a sign of nonrandom correlations in the matrix, and for this reason we refer to these principal components as signal. A reactivation time series (**Fig. 6c**) measured the instantaneous match of this cell assembly to the ongoing activity[35].

Statistical analysis. We performed one-way ANOVA with multiple comparisons wherever significance assessment was required (**Figs. 2b**, **3c**, **4c** and **5b**, and **Supplementary Figs. 1d**,**f**,**g**, **2f**, **4a**,**b** and **8**). We also used a paired t test for comparisons between early and late trials and features of SWS_{pre} and SWS_{post} and other relevant comparisons (for example, time to reward and proportion of unsuccessful trials in and early and late trials, and time to reward comparison between $block_1$ and $block_2$ for second sleep data set; delta wave density, PSD amplitude changes, duration changes, etc., in SWS_{pre} and SWS_{post}; change in firing rate of units, reactivation strengths in post awake and SWS epochs, etc.). Unpaired t test were also used for several comparisons namely, proportion of cells with significant MD_{change} in direct and indirect unit pools, STA_{amp} change for units recorded with silicon probes versus multielectrode arrays, delta-triggered reactivation strength comparison in SWS_{post} from robust and poor learning sessions, etc. We also used linear regression to evaluate trends between firing rate change of individual units versus STA amplitude and PSD amplitude versus STA amplitude. Kolmogorov–Smirnov test was used for comparing burst distribution in interspike interval distributions from SWS_{pre} and SWS_{post}. We used the spearman correlation between time spent in SWS and improvements in task performance because the distribution of SWS durations were not normal. For similar reasons, we used spearman correlation between the firing rate modulation and STA effect size. Logrank test was used to determine probability of higher strength reactivation strength in SWS_{post} after robust or poor learning and also in SWS_{post} versus post awake epochs after robust learning. If there was higher likelihood of greater reactivation strength in comparison to SWS_{pre}, the difference of counts of reactivations strengths of higher magnitude would show right-skewed histogram and logrank test can test if such right-skewed curves are different from each other.

A **Supplementary Methods Checklist** is available.

46. Hill, D.N., Mehta, S.B. & Kleinfeld, D. Quality metrics to accompany spike sorting of extracellular signals. *J. Neurosci.* **31**, 8699–8705 (2011).
47. Fee, M.S., Mitra, P.P. & Kleinfeld, D. Automatic sorting of multiple unit neuronal signals in the presence of anisotropic and non–Gaussian variability. *J. Neurosci. Methods* **69**, 175–188 (1996).
48. Bjorvatn, B., Fagerland, S. & Ursin, R. EEG power densities (0.5–20 Hz) in different sleep-wake stages in rats. *Physiol. Behav.* **63**, 413–417 (1998).
49. Phillips, K.G. *et al.* Decoupling of sleep-dependent cortical and hippocampal interactions in a neurodevelopmental model of schizophrenia. *Neuron* **76**, 526–533 (2012).
50. Pack, A.I. *et al.* Novel method for high-throughput phenotyping of sleep in mice. *Physiol. Genomics* **28**, 232–238 (2007).
51. Mitra, P. & Bokil, H. *Observed Brain Dynamics* (Oxford University Press, 2008).

doi:10.1038/nn.3759