Noradrenergic locus coeruleus ensembles evoke different states in rat prefrontal cortex

Shahryar Noei [1,2], Ioannis S. Zouridis [3,6], Nikos K. Logothetis [6,7], Stefano Panzeri [1] **, Nelson K. Totah [4,5,6] **

[1] Neural Computation Laboratory, Istituto Italiano di Tecnologia, Rovereto, Italy

[2] Center for Mind/Brain Sciences, University of Trento, Rovereto, Italy

[3] Graduate Training Centre of Neuroscience, University of Tübingen, Tübingen, Germany

[4] Helsinki Institute of Life Science (HiLIFE), University of Helsinki, Helsinki, Finland

[5] Faculty of Pharmacy, University of Helsinki, Helsinki, Finland

[6] Dept. of Physiol. of Cog. Processes, MPI Biological Cybernetics, Tübingen, Germany

[7] Division of Imaging Sci. and Biomed. Eng., University of Manchester, Manchester, UK

** Senior authors

Abstract

Identifying the neurons that control global brain state has been a fundamental topic of research that has largely focused on diffusely-projecting neuromodulatory centers, such as the locus coeruleus (LC). This noradrenergic brain stem nucleus, which projects throughout the forebrain, is thought to act as an “undifferentiated state controller” across all forebrain targets because LC neurons spike synchronously. However, recent work demonstrated ensembles in the LC and therefore made targeted neuromodulation a possibility. In order to demonstrate that LC ensembles cause targeted neuromodulation, it is necessary to resolve LC ensemble dynamics over time in relation to ongoing cortical states. Here, we used non-negative matrix factorization on LC single unit recordings to investigate the spatial and temporal properties of ensemble activation patterns. We assessed the potential for targeted neuromodulation of the prefrontal cortex (PFC) using LC ensemble activity-triggered local field potential (LFP) power spectrograms. We analyzed 285 single units recorded from 15 urethane-anesthetized rats (range of 5 to 34 simultaneously recorded units). LC ensembles became active at different times. Analysis of auto-correlograms and ensemble-pair cross-correlograms demonstrated that self- and lateral-inhibition of activity is a property of LC ensembles, which may contribute to their sparse activity. Neuromodulatory effects on cortical state were diverse across ensembles. We observed four types of ensemble-triggered LFP spectrograms in the PFC. These results demonstrate that the LC is capable of differentiated neuromodulation of its forebrain targets by dynamic firing patterns across subsets of LC neurons.

Introduction

Internal, spontaneously occurring brain states are associated with changes in wakefulness, perceptual ability, and reaction times (1–3). Behavioral transitions, such as waking from sleep or reacting more quickly to stimuli during stress, involve changes in global brain state. It remains unclear exactly which neurons could control brain state globally.

Brainstem neuromodulatory centers, such as the noradrenergic locus coeruleus (LC), are a likely global state-controller (1, 2, 4). The LC projects globally throughout the central nervous system and releases norepinephrine to modulate neuronal excitability (5–11). Global noradrenergic neuromodulation can result in behavioral transitions, such as: awakening from sleep or anesthesia, altering locomotion patterns (increased generalized movements and decreased reaction times), and improving perceptual sensitivity and attentional focus (12–24). The brain state change that occurs spontaneously during behavioral transitions, such as awakening, also resembles the brain state evoked by stimulation of the LC in awake or anesthetized animals (9, 16, 24, 25). Thus, the LC is thought to be at least one group of neurons that can alter global brain states associated with behavioral state transitions.

LC neurons do not necessarily provide a global neuromodulatory signal, however. Recent experiments that recorded from large populations of well-isolated LC single units demonstrated that LC single units spike together in ensembles (26). If the units within an ensemble all project to one (or a few) brain region(s), the neuromodulation of brain state by the LC could be targeted or "semi-global." Studies characterizing the projections of LC neurons have shown that some LC neurons have highly localized (single target) or semi-global projections (26–30). Evidence for LC ensemble activity and targeted projections, together, support the idea that the LC may be capable of controlling brain state locally; however, the relationship between brain state and LC ensemble activity has not been studied.

This relationship cannot be studied with the method used previously to detect LC ensembles because it provides only a static "snapshot" of which single units spike together on average throughout an experiment (26, 31). Here, we apply a different tool (non-negative matrix factorization, NMF (32)), which detects LC ensembles and resolves their temporal dynamics over time. We studied the relationship between local brain state and LC ensemble activity by triggering

cortical local field potentials (LFPs) on the activation times of different LC ensembles. We applied this method to a prior data set (26) and a new data set. The recordings were made under urethane anesthesia, an ideal model for studying the transition between brain states (33). We found that LC ensembles detected with NMF were similar to those detected using the static "snapshot" approach (i.e., the graph theoretic analysis of pairwise correlations used in prior work (26)). Specifically, ensembles were sparse, lacked a spatial topography in the LC, and were LC cell type-specific. With the availability of ensemble temporal dynamics, we were able to study interactions among LC ensembles. We show that many of the principles that apply to LC single unit firing (such as self-inhibition and lateral-inhibition (34, 35)) occur at the level of collections of co-active LC single units. However, these effects are less prominent and allow coordinated co-activation of multiple LC ensembles. These diverse and independent patterns of activity among LC ensembles may allow targeted neuromodulation. By examining ongoing cortical LFPs, we found that cortical state depended upon which ensemble was active and did not fit the canonical model of LC activation causing a global and homogenous brain state change.

Methods

Data collection: Recording and signal acquisition

Experiments were carried out with approval from the local authorities and in compliance with the German Law for the Protection of Animals in experimental research (Tierschutzversuchstierverordnung) and the European Community Guidelines for the Care and Use of Laboratory Animals (EU Directive 2010/63/EU). Male Sprague-Dawley rats (350 - 450 g) were used (specific pathogen free, Charles River Laboratories, Sulzfeld, Germany). They were pair housed. Experiments were carried out on a 08:00 lights off / 20:00 lights on cycle. A sub-set of the data were collected from rats used in a prior study (26).

Rats were anesthetized using an intra-peritoneal (i.p.) injection of urethane at a dose of 1.5 g/kg body weight (Sigma-Aldrich, U2500). Surgical procedures were as described in prior work (26).

Electrodes targeted the LC and the prelimbic division of the medial prefrontal cortex. The coordinates for LC were 4.0 mm posterior from lambda, 1.2 mm lateral from lambda, and approximately 6.0 mm ventral from the brain surface (implanted at a 15 deg posterior angle). The coordinates for the cortex were 3.0 mm anterior and 0.8 mm lateral from bregma and 3.0 mm ventral from the brain surface. The LC electrode was targeted based on standard

93 electrophysiological criteria (see prior work for a detailed description (26)). At the end of the
94 recording, we administered clonidine (0.05 mg/kg) i.p. (Sigma-Aldrich, product identification:
95 C7897) to confirm cessation of noradrenergic neuron spiking. We also verified LC targeting in
96 most experiments using histological examination of coronal sections (50 um thick) that were
97 stained for Cresyl violet or a DAB and horse radish peroxidase reaction with hydrogen peroxide
98 to visualize an antibody against tyrosine hydroxylase.

99 *Data collection: electrophysiology*

100 The LC was recorded using a multi-channel silicone probe (NeuroNexus, Model: A1x32-Poly3-
101 10mm-25s-177-A32). The impedance of the electrodes was ~1.0 to 2.0 MOhm. Cortical local field
102 potentials were recorded using a single tungsten electrode with an impedance of 200 – 800 kOhm
103 (FHC). A silver wire inserted into the neck muscle was used as a ground. Electrodes were
104 connected to a pre-amplifier (in-house constructed) via low noise cables. Analog signals were
105 amplified (by 2000 for LC and 500 for cortex) and filtered (8 kHz low pass, DC high pass) using
106 an Alpha-Omega multi-channel processor (Alpha-Omega, Model: MPC Plus). Signals were then
107 digitized at 24 kHz using a data acquisition device (CED, Model: Power1401mkII).

108 *Ensemble Detection*

109 To search for common patterns of firing among Locus Coeruleus cells, we utilized non-negative
110 matrix factorization (NMF (32)) decomposition of the single unit by time matrix of spike rates. This
111 method linearly decomposes the data into a set of non-negative basis functions (modules) using
112 non-negative coefficients. The basis functions can be interpreted as the basic components that
113 make the raw data (32, 36). The non-negativity constraint improves the interpretability of the basis
114 functions and the coefficients, as they can be assumed to be a set of temporal patterns of positive
115 firing rates and the recruitment of single units in that pattern (36).

116 Here, we used “space only non-negative matrix factorization”, which is a specific design of NMF
117 that is especially suitable for decomposing the matrix of firing rates into common patterns across
118 subsets of single units (36):

119 $$\mathrm{R} = \mathrm{WH} + \text{residuals}$$

120 $\mathrm{R} \in \mathbb{Z}_{+}^{T \times N}$ is the spike trains matrix of N single units binned every t time, resulting in total T bins.
121 $\mathrm{W} \in \mathbb{R}_{+}^{T \times K}$ is the basis function consisting of K modules and $\mathrm{H} \in \mathbb{R}_{+}^{K \times N}$ is the coefficient matrix.

The residuals are the unexplained random white noise. For decomposition, we used the multiplicative update rule to minimize the Frobenius norm between the original and the reconstructed data (32).

Using this formula, we can assume each column of W as a temporal module representing underlying firing pattern in the neural population and each row of H as the amount of the engagement of each single unit in that particular firing pattern. In other words, coefficients in the rows of H represent the amount of engagement of each single unit to a specific pattern of firing over time.

Two hyperparameters were selected manually: the binning resolution t and the number of modules k. We binned the spike trains in 100 msec based on prior work demonstrating that pairs of Locus Coeruleus single units are predominantly synchronized on a timescale of approximately 100 msec or less (26). We also binned the spike trains with a finer temporal resolution (20 and 10 msec), but the result of NMF was ensembles consisting of only a single unit (which is, by definition, not an ensemble). On the other hand, a longer temporal resolution led to excessive synchrony (i.e., one large ensemble containing all single units) and, thus, failed to detect real ensembles.

We estimated the ”optimal” number of modules (i.e., ensembles) for each rat by varying the number of modules from one up to the total number of single units and looking at the amount of the variance explained by the decomposition. We chose the number of modules based on the first elbow of this distribution after which at least 60% of the original data variance is explained. These criteria provided an estimate of the optimal number of modules in the recorded data, but may not represent the real number of ensembles in the LC due to lack of access to all the single units in the LC. After determining the number of modules, the final decomposition was performed 5 times for each rat to avoid falling into local minima. The final decomposition was chosen as the one leading to the minimum error. One concern with assembling methods, like NMF, is that it can return modules that consist of random single units that do not actually fire together. We ensured that the detected ensembles were repeatable (i.e., non-random) by checking the repeatability and degree of clustering in comparison to random assignment of single units to ensembles. We enforced a “hard clustering” procedure, which forces each single unit to fall into one and only one ensemble. This was done by dividing each column of H by its maximum and removing the values below 1. Hard clustering was done for the sole purpose of checking the repeatability of the method across 5 decompositions (for each rat) and measuring their agreeability using the Rand Index (37). We compared the average of the Rand index for each animal with 100 repetitions of 5

random clustering. The average Rand Index from the hard clustering was always greater than the top 5% of the distribution of mean Rand Indices resulting from random clustering. Therefore, NMF decomposition produced meaningful and repeatable ensembles.

However, hard clustering for detecting activation of ensembles may not be biologically plausible because single units may spike as members of different ensembles at different times. Therefore, the analyses in this paper did not use hard clustering and instead used a threshold-crossing to define when ensembles were active and which neurons were active in the ensemble. In order to detect ensemble activations based on the NMF coefficients, we first normalized the columns of H to the minimum and maximum and then set a threshold based the distribution of coefficients. The distributions had three peaks: one with excessively small values (effectively, no activity of the single units in the ensemble), one with intermediate (values), and one with larger values. Participation of single units in an ensemble was defined as activation coefficients from H crossing the 95_{th} percentile threshold of the distribution of coefficient values for that rat. This threshold effectively selected the values after the first peak (no participation of the single units) in every case. Coefficients below this value were set to zero and values above the threshold were set to one. In the resulting binary version of the matrix, H, a value of 1 represented spatial modules corresponding to a single unit belonging to an ensemble.

The columns of the W matrix correspond to a set of temporal bases that represent the firing patterns underlying each ensemble. We divided these continuous non-negative values into binary values using the same method explained above for the spatial modules. The binary version of the matrix, W, (hereafter referred to as “activation coefficient matrix”) was used to determine whether an ensemble is active or not for each time bin.

Evaluating physical clustering of ensembles in space

To assess whether single units within an ensemble tend to cluster in the LC, we measured the pairwise distance between the units within each ensemble. The location of each unit was assigned to the electrodes on which the maximal waveform was recorded. After assigning the unit location, Euclidian pairwise distances of the units inside each ensemble was calculated and the distributions were plotted.

Assessing the proportion of single unit types in the ensembles

We observed that single units of the same type were more likely to fire as an ensemble. We determined if this was more likely than by chance by computing the exact probability of having ensembles of the same single unit type. The probabilities were computed by the means of repetition of random sampling (assembling) without replacement. The number of units in the sample was fixed to the number of single units in the ensemble. The number of repetitions for each rat was the number of ensembles that consisted of only of one type of single unit.

Analysis of interactions between pairs of ensembles

The interactions between the ensembles was measured using cross-correlograms. Cross-correlograms were calculated in a window of 2000 msec with a bin size of 100 msec. The cross-correlograms were compared to 1000 surrogate cross-correlograms by jittering the activation times uniformly between ±1000 msec (38). Significant excitatory (or inhibitory) interactions were those that had cross-correlogram bins which crossed the upper (or lower) 1% of pairwise coincidental activations observed in the surrogate data.

We checked for the degree of synchrony between the ensembles that had a significant excitatory interaction at time 0. We introduced a synchrony index as follows:

$$synch = \left(\frac{2 * c_{ij}}{\tau_i + \tau_j}\right) * 100$$

Where c_{ij} is the number of times the two ensembles are coactive and τ_i, τ_j are the number of active times for each ensemble. We calculated this synchrony index for all the ensembles that have significant cross-correlogram at time 0 and plotted the histogram.

Analysis of self-inhibition by ensembles

Self-interactions were measured using auto-correlograms. These were calculated in a 1000 msec time window using a 100 msec time bin. The significance of inhibition was assessed by comparing the observed auto-correlogram versus 1000 surrogate auto-correlograms. The procedure was identical to that explained above for the cross-correlograms.

Peri-event time histogram calculation and clustering

To calculate the Peri-Event Time Histograms (PETHs), the spike times of single units inside and outside of an ensemble were aligned to events (at t = 0 msec), which were the ensemble activation times. We examined activity during a window from 100 msec before up to 400 msec after the ensemble activation times and used 1 msec bins. For each single unit, we calculated the average spike rate across activation events as though they were “trials”. These PETHs were smoothed with a Gaussian kernel (10 msec width). The PETH for each ensemble was obtained by averaging PETHs across all single units within the ensemble.

PETH clustering was done in two steps. First, the dimensionality of the original PETHs in time was reduced using the Principle Component Analyses (PCA) (39). We observed that reducing the dimension to 2 would explain more than 95% of the variance in the original data. After visualizing the data in the two dimensions we observed 3 non-circular masses of data, so we decided to cluster the data in 3 groups using a Gaussian Mixture Model (GMM) (40) with 3 repetitions and full covariances.

Burst rate calculation

We studied bursting in two ways. The first method merged the spike times of single units within each ensemble. The purpose of this analysis was to study the temporal structure of spiking across all units within the ensemble. The second method examined the bursting of each single unit within each ensemble. In both cases, burst rate was also calculated for all single units outside of the ensemble for comparison. We defined a burst as each occurrence of 2 or more consecutive spikes with an inter-spike interval of less than 80 msec. We measured burst rate by dividing the number of bursts happening during the total time the ensemble was active. When comparing with the burst rate of single units outside of the ensemble, we used the ensemble inactive periods.

LFP spectrogram modulation index

We investigated the relation between the activation of LC ensembles and cortical LFP power by triggering LFP spectrograms on ensemble activation events. Spectra were computed using multitaper method implemented in Chronux toolbox (41) with 3 tapers and time bandwidth product of 5. Short-time Fourier transforms were computed in a 10 msec moving window with a duration of 200 msec. The resulting spectral resolution was ~4 Hz and the temporal resolution was 10 msec. We computed an ensemble activation-triggered modulation index for each spectrogram in

order to characterize the effects of LC ensemble activation on the cortical LFP power spectrum. We averaged the spectrogram in time at each frequency for the baseline duration (400 msec before the ensemble being active), then we subtracted the baseline averaged spectrogram from the original spectrogram at each time step and divided by their sum:

$$MI_s(t,f) = \frac{S(t,f) - S_{baseline}(f)}{S(t,f) + S_{baseline}(f)}$$

This value varies between -1 to 1 and can potentially describe the change in the cortical LFP power around ensembles activation.

Spectrogram clustering

To assess the diversity of LC ensemble activation-triggered cortical responses, the spectrogram modulation indices were clustered. Spectral images were vectorized and k-means clustered (42). We clustered the modulation indices into 4 different clusters setting the correlation as the distance function of the algorithm. We choose these values by first varying the number of clusters from 1 to 22 (approximately one-quarter of the data) using for different distance functions (Euclidian, cosine, cityblock and correlation). Then, we looked at the normalized error (error divided by the maximum error) curve and chose the elbow as the point where the drop in the error is less than 5% for each. After fixing the number of clusters, we performed a sillohute analysis (43) or different distance functions to check for the consistency of the clustering. Together, by comparing the average silhouete value, the normalized error, and equality of samples in the clusters, we found that using correlation as the distance function with 4 clusters leads to acceptable results.

We performed the clustering and visualized the average spectrogram modulation indices for each cluster and asssessed the significance at each time and frequency by pooling all the ensembles spectrogram modulation indices inside each cluster and compared the median of the population against zero using Wilcoxon signed rank test (5% significance level). The p-values were corrected for multiple comparasions using Benjamini's & Hochberg's method for false discovery rate (44).

Detecting ensembles preferred cortical state

To measure whether different ensembles are preferentially active in specific cortical states defined by very slow LFP oscillations (peak <1Hz), slow LFP oscillations (peak between 1-2 Hz), an "activated state" of increased high frequency LFP oscillations (> 20 Hz) and decreased low

frequency LFP oscillations (< 2 Hz), or a mixture of slow and activated states, we first defined cortical state in windows (7.5 sec duration). The distribution of LFP voltages was obtained for each window. The distribution was tested for bimodality using Hartigan's Dip Test ($p<0.05$). A significant dip test selected epochs that were bimodal and therefore either very slow oscillations or slow oscillations. We also separated those states with a significant dip test into very slow oscillation states and slow oscillation states using the proportion of the power spectrum of each LFP epoch that was very low frequency (<0.4 Hz). The distribution of power ratios was bimodal, which suggested that epochs of LFP clustered into very slow oscillation and slow oscillation states. A non-significant dip test selected for epochs of LFP that were relatively flat (activated state or mixture of activated and slow oscillations). We separated activated states from mixture states by examining the kurtosis of the LFP voltage distribution, with high kurtosis values indicating a sharply peaked distribution with very little variability (activated state). Each 7.5 second epoch of LFP (and its voltage distribution) was thus associated with 3 values: a dip test p value, kurtosis, and power ratio. These values were used with K-means clustering to assign each LFP epoch a state: activated, mixture (activated and slow oscillations), slow oscillations, very slow oscillations, and unclassified.

Finally, we characterized each ensemble's preference for being active during a specific cortical state. We computed the likelihood of an ensemble being active at a given state (using Bayes rule) and compared it to 1000 surrogate likelihoods computed by shuffling activation times, which keeps the number of activation times constant). If the likelihood at a certain state crosses the 95_{th} percentile of the surrogate distribution, we assumed that the ensemble has a preference of being active during that state. If the likelihood for an ensemble crossed the statistical threshold for more than one state, we marked the preferred state as the one that had the larger likelihood value.

Results

Ensemble detection using non-negative matrix factorization

The ensembles were detected using NMF on the single units recorded from each rat. This method decomposes the single unit time binned spike counts into two non-negative matrices; one consists of a set of basis firing patterns over time and the other shows how much each single unit participates in that particular firing pattern (see Methods section on Ensemble detection for further details). **Figure 1A** shows an example of how NMF works on hypothetical single unit spiking data. Based on prior studies on LC neurophysiology, we chose 100 msec as our binning length and

296 applied NMF on LC single unit data (26). We estimated the ”optimal” number of ensembles for
297 each rat by varying the number of modules from one up to the total number of single units and
298 looking at the amount of the variance explained by the decomposition (see methods section for
299 additional details). We chose the number of modules based on the first elbow of this distribution
300 after which at least 60% of the original data variance is explained (**Supplementary Figure 1**). We
301 found a total of 146 ensembles from 285 single units. Note that a single unit can spike in more
302 than one ensemble. **Figure 1B** shows two exemplar ensembles (top and bottom). The left panels
303 depict the spike rasters of single units which belong to the ensemble and the right panels show
304 spike rasters of single units outside the ensemble. Spike rasters are aligned to the ensemble
305 activation times. These examples clearly show that single units, which are assigned to the same
306 ensemble by NMF, increase their firing rate during ensemble active times (0 to 100 msec around
307 ensemble activation). On the other hand, units not assigned to the ensemble maintain their
308 previous pattern of activity. Decomposing the pooled spike rates of single units to ensembles can
309 be better conceptualized by looking at **Figure 1C**. Here, the pooled spike rate of single units
310 during a selected window of recording from an example rat has been decomposed into five
311 ensembles by NMF. Although multiple ensembles may sometimes activate synchronously, the
312 example shows that they can have very sparse and different firing patterns in general.

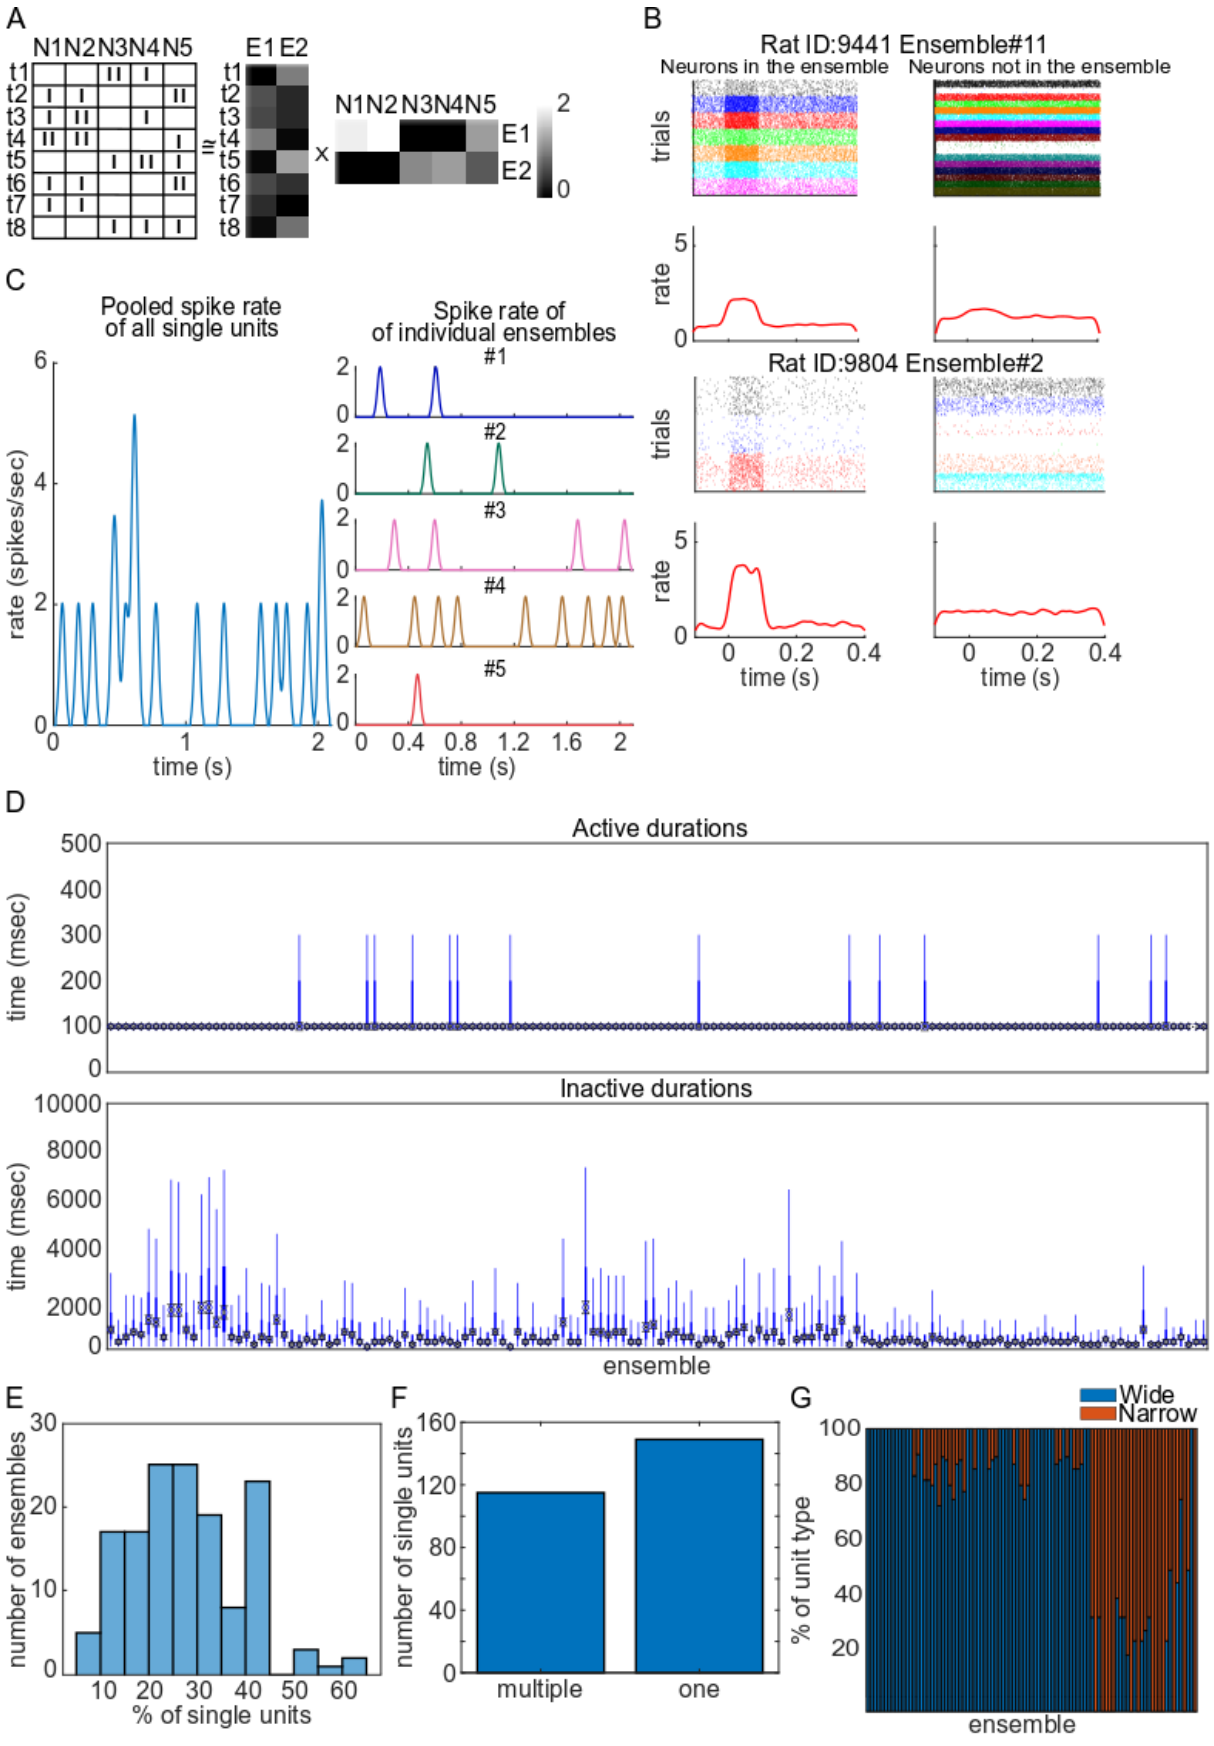

313

Figure 1. Ensemble detection using NMF decomposes LC population spiking into ensembles that have heterogeneous spatio-temporal properties. (A) An example of how NMF works on hypothetical data from 5 single units. The spike times are shown as raster ticks for each unit (N1 to N5) over time (y-axis). Time has been binned (t1 to t8) so that spikes per bin can be counted. The matrix of spike counts per time bin (rows) over units (columns) can be decomposed into the two base matrices illustrated on the right with values (greyscale) corresponding to the amount of spiking (middle panel) and how much each single unit participates in that particular spiking pattern (right panel). Visual inspection of the spike rasters (left panel) shows that N1, N2, and N5 tend to spike in the same time bins (t2, t3, t4, t6, t7). Therefore, those units are an ensemble, which is active at those time points. This ensemble, labeled E1 in the plots on the middle and right, has lighter values (stronger activation coefficients) at times t2, t3, t4, t6, and t7 (middle plot) with higher values (lighter color) for higher population spike count within this ensemble. The plot on the right for E1 shows lighter colors (stronger activation coefficients) for N1, N2, and N5 which indicates that they fire together. Note that the activation coefficients are higher for N1 and N2, which also tend to respond together more regularly in the spike rather. **(B)** The spike rasters and PETHs are shown for two exemplar LC ensembles recorded from two rats. The left panels show spike rasters of the single units inside the ensemble aligned to the ensemble activation time (t = 0 sec). Below the rasters, the PETHs are shown. The right panels depict the ensemble activation-triggered spiking of single units that were *not assigned* to that ensemble by NMF. The plots show that units inside the ensemble increased their firing rate at ensemble activation times, whereas units not assigned to the ensemble did not change their firing rate. **(C)** The pooled spiking activity of the all single units in one example rat (left panel) decomposed into 5 different ensembles. The right panel shows the activation coefficients over time for each of the 5 ensembles. As can be seen by visual inspection, summing the ensemble activity patterns results in an approximate reconstruction of the original population spiking on the left. Therefore, NMF is able to deconstruct population spiking patterns into sub-sets of co-active single units. **(D)** The box plots show the distributions of how long ensembles were consecutively active (top panel) and inactive (bottom panel). Each boxplot illustrates the distribution for one ensemble. Ensembles tend to be active for only 100 msec, but can be inactive for a wide variety of durations which yields heterogenous activation of different ensembles at different times. **(E)** The histogram shows the distribution of ensemble size, in terms of the percentage of simultaneously recorded single units that were assigned to an ensemble. On average, each ensemble consisted of 27% of the single units recorded in that experiment. **(F)** This histogram shows the number of single units that fall into a single ensemble or multiple ensembles. There was a preference for units to participate in only one ensemble. **(G)** The percent of each unit type (wide or narrow spike waveform) making up each ensemble is plotted across ensembles (x-axis). Ensembles are either a single unit type or consist of mostly a single unit type.

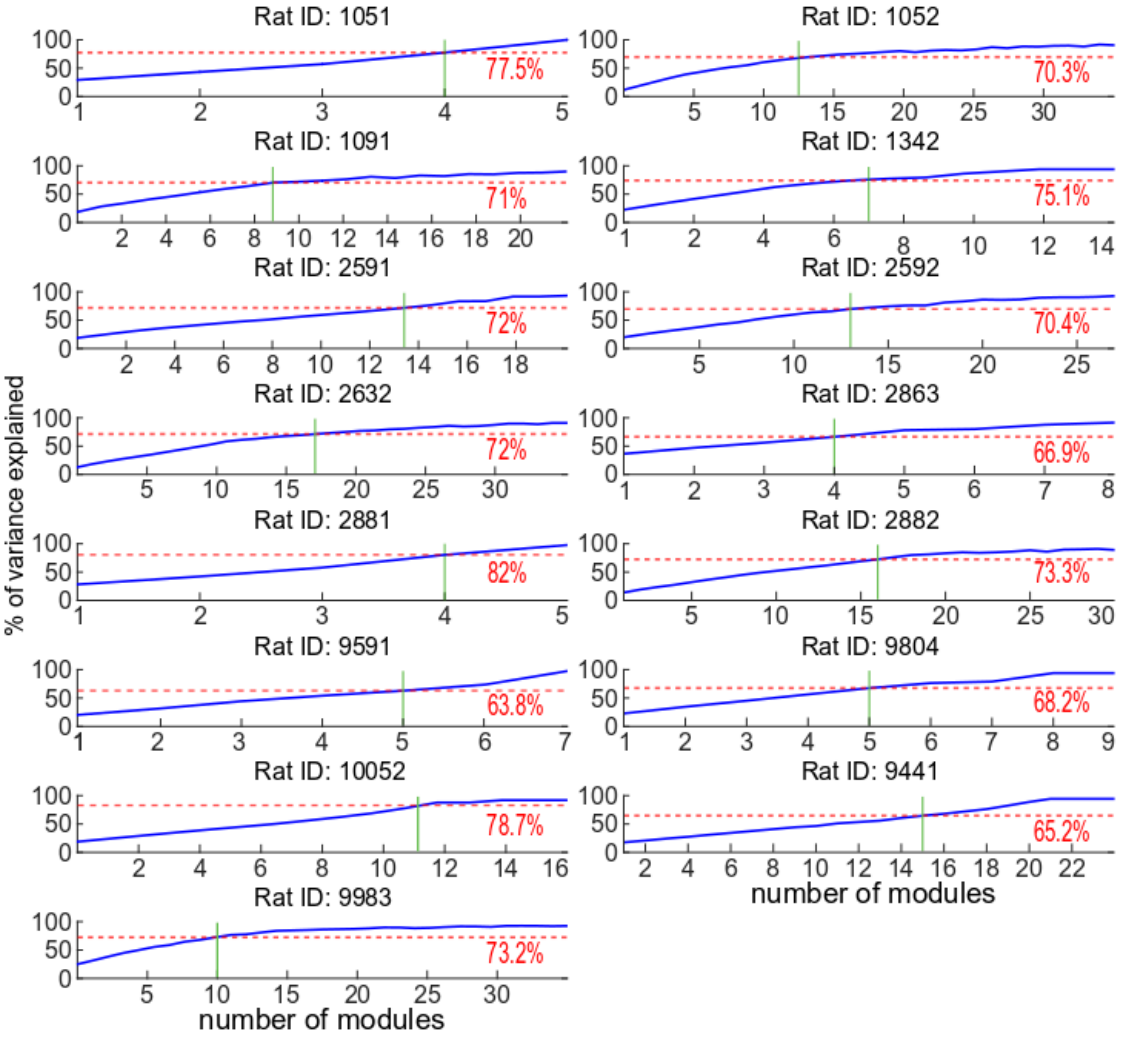

Supplementary Figure 1. Data underlying the choice of the optimal number of ensembles in each experiment. Each panel depicts the percentage of explained variance using NMF versus the number of the modules (ensembles) for each rat. Solid green lines show the number of selected modules based on the criteria of first elbow after at least 60% of variance is explained. The dotted red lines show the amount of the explained variance at the selected number of modules.

The spatio-temporal scale of ensemble activation is heterogeneous

We investigated the duration that ensembles were active or inactive. **Figure 1D** presents evidence that most ensembles did not stay active for longer than 100 msec. However, the inactive duration varied across ensembles. Our analysis of the temporal dynamics of ensemble activity demonstrates that ensembles are "quiet" for a wide variety of durations, before being briefly active for ~100 msec. Therefore, at the overall LC population level, various ensembles briefly respond at diverse times, which yields a temporally-heterogeneous population code.

We next studied the spatial characteristics of the ensembles. We first assessed the size of each ensemble. In **Figure 1E**, we report the relative number of neurons in each ensemble relative to the total number of the neurons recorded from that rat. We found that, on average, 27% of single units were active in ensembles (although some units could participate in more than one ensemble, as discussed in the next paragraph). Ensembles could vary in size, which ranged from 6% to 62% of the recorded single units firing in an ensemble (**Supplementary Figure 2A**). These single units were spread throughout the LC with no topographical organization (**Supplementary Figure 2B**). Among 146 ensembles, only 23 ensembles had median pairwise distribution smaller than 50 um. The diffuse spatial arrangement of single units within ensembles detected with NMF agrees with prior work on LC ensembles that used graph theoretic analysis of pairwise spike count correlations to detect ensembles (26).

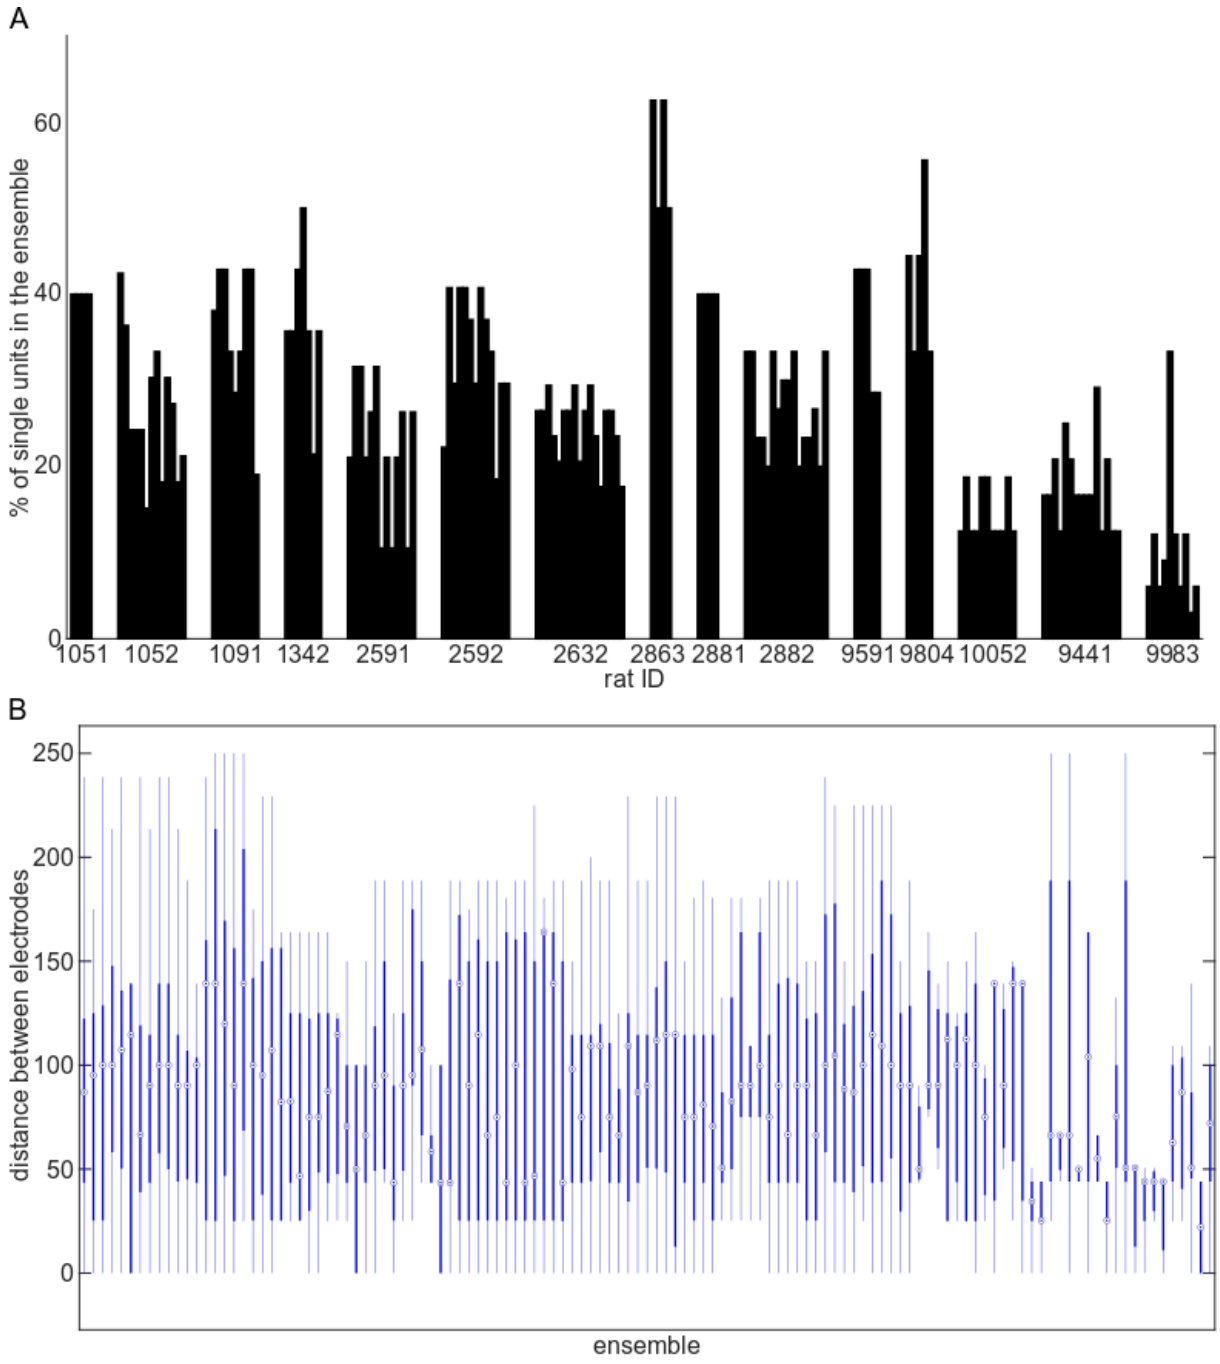

Supplementary Figure 2. The spatial properties of the detected ensembles. **(A)** Bar plot showing the percentage of all simultaneously recorded single units within each ensemble. The percentage was calculated as the number of single units inside the ensemble divided by the total number of single units recorded for that rat. Each ensemble is a bar. The bars are grouped by rat. Note that a single unit can be part of more than one ensemble. Overall, the results suggest that the ensembles can vary in size. **(B)** Boxplots showing the pairwise Euclidian distance among the single units inside an ensemble. Ensembles with only two single units were excluded from this plot. The distributions indicate that ensembles are spatially diffuse.

We investigated how many of the single units were assigned to a single ensemble, multiple ensembles, or none. Out of 285 single units, 115 single units spiked as part of multiple ensembles (40.4%), 149 spiked in only a single ensemble (52.3%), and the remaining 21 units did not form any ensemble with the other single units (**Figure 1F**). We tested the hypothesis that a larger proportion of units fired exclusively in only one ensemble as opposed to participating in multiple ensembles. Single units tended to have a "preferred" single ensemble (binomial test, $p = 0.0421$).

LC ensembles were unit type-specific

We next examined whether LC single units of the same type tended to spike in the same ensemble or not. These units are distinguishable by their extracellular waveform shape and are therefore termed "narrow" or "wide" type units (26). The recent work that first identified narrow and wide LC unit types showed that LC units of the same type tended to form ensembles (26). This prior work used graph theoretic measures to detect ensembles which, in contrast to the NMF method used here for ensemble detection, does not allow one to assess ensemble activation times or participation of units in more than one ensemble. The present analysis was therefore important for determining if NMF, a mathematically completely unrelated method used for ensemble detection, would also provide evidence of LC unit type-specific ensembles. **Figure 1G** reports the percent of each unit type participating in each ensemble (after removing rats in which only a single unit type was recorded). Visual inspection of the plot clearly shows that units of same type tend to form ensembles. We assessed if these proportions were statistically different from what would be expected by chance using random sampling (see methods). Our results show that, for all rats, the probability is below 0.05, thus suggesting that this number of unit type-specific ensembles is not due to chance.

LC ensemble activation is associated with burst firing

We examined the characteristics of spiking during LC ensemble activation by examining the single unit spike rates and burst rates of single units within an ensemble versus those not in the ensemble. We characterized the average spike rate of units within the ensemble when the ensemble was active by calculating peri-event time histograms (PETH) of spike rate for each ensemble. Each event was an ensemble activation time. Single unit spike times were aligned to the ensemble activation times from 100 msec before each ensemble activation event until 400 msec after it. In order to assess whether the spike rate differed across ensembles, we clustered

the PETHs (one for each of the 146 ensembles) using PCA and GMM (see methods). After visualizing the data in two dimensions we observed 3 non-circular masses of data (**Supplementary Figure 3A**). therefore, we divided the PETHs into 3 groups, which were associated with low, medium, and high changes in spike rate, but similar activation durations (**Figure 2A**). Most ensembles were characterized by a low or medium change in spike rate corresponding to an increase of 1 to 3 spikes per sec (**Figure 2B**). In the maximal case, single unit spike rate increased by 7 spikes per sec (**Figure 2A**, magenta line), but this was the smallest group of ensembles (**Figure 2B**, magenta). We also assessed the size of the ensembles according to their activation size (low, medium, and high PETH clusters). The median of number of single units that fell into ensembles belonging to each cluster of PETHs decreased from cluster 1 to cluster 3 (one-sided Wilcoxon rank sum test, 5% significance level, corrected for multiple comparisons with the Bonferroni-Holm correction (45)). Therefore, activation of different ensembles produced different degrees of spiking output that decreased with the number of units in the ensemble. This is probably due to the fact that, as the ensemble size grows larger in terms of the number of units, there will be less synchrony among all of the units in terms of precise spike timing and not all the units within the ensemble are increasing their firing rate on each activation instance of the ensemble.

We assessed the diversity across ensembles of another form of LC output, which is the tendency to fire in bursts. We defined a burst as each occurrence of 2 or more consecutive spikes with an inter-spike interval of less than 80 msec. We chose 80 msec by looking at the pooled histogram of inter-spike intervals (ISI) of all the recorded single units (**Supplementary Figure 3B**). The 5_{th} percentile of the ISI distribution is 78 msec. This 80 msec threshold for bursting is consistent with the physiology of LC neurons in which spikes with an ISI of less than 100 msec are more likely to propagate and release norepinephrine (46). Single units within an ensemble, when merged into a single population spike train, burst more often when the ensemble was active than when the ensemble was in an inactive state (**Figure 2C**, two-sided Wilcoxon rank sum test, $Z = 13.128$, $D = 1.011$, power = 0.99, $p < 0.0001$). We also calculated the burst rate without merging the spike times across single units. Within ensembles, single units tended to burst more during the ensemble active times than during inactive times (**Figure 2D**, two-sided Wilcoxon rank sum test, $Z = 15.217$, $D = 0.457$, power = 0.99, $p < 0.0001$). These results demonstrate than LC ensemble activation is associated with increased burst firing.

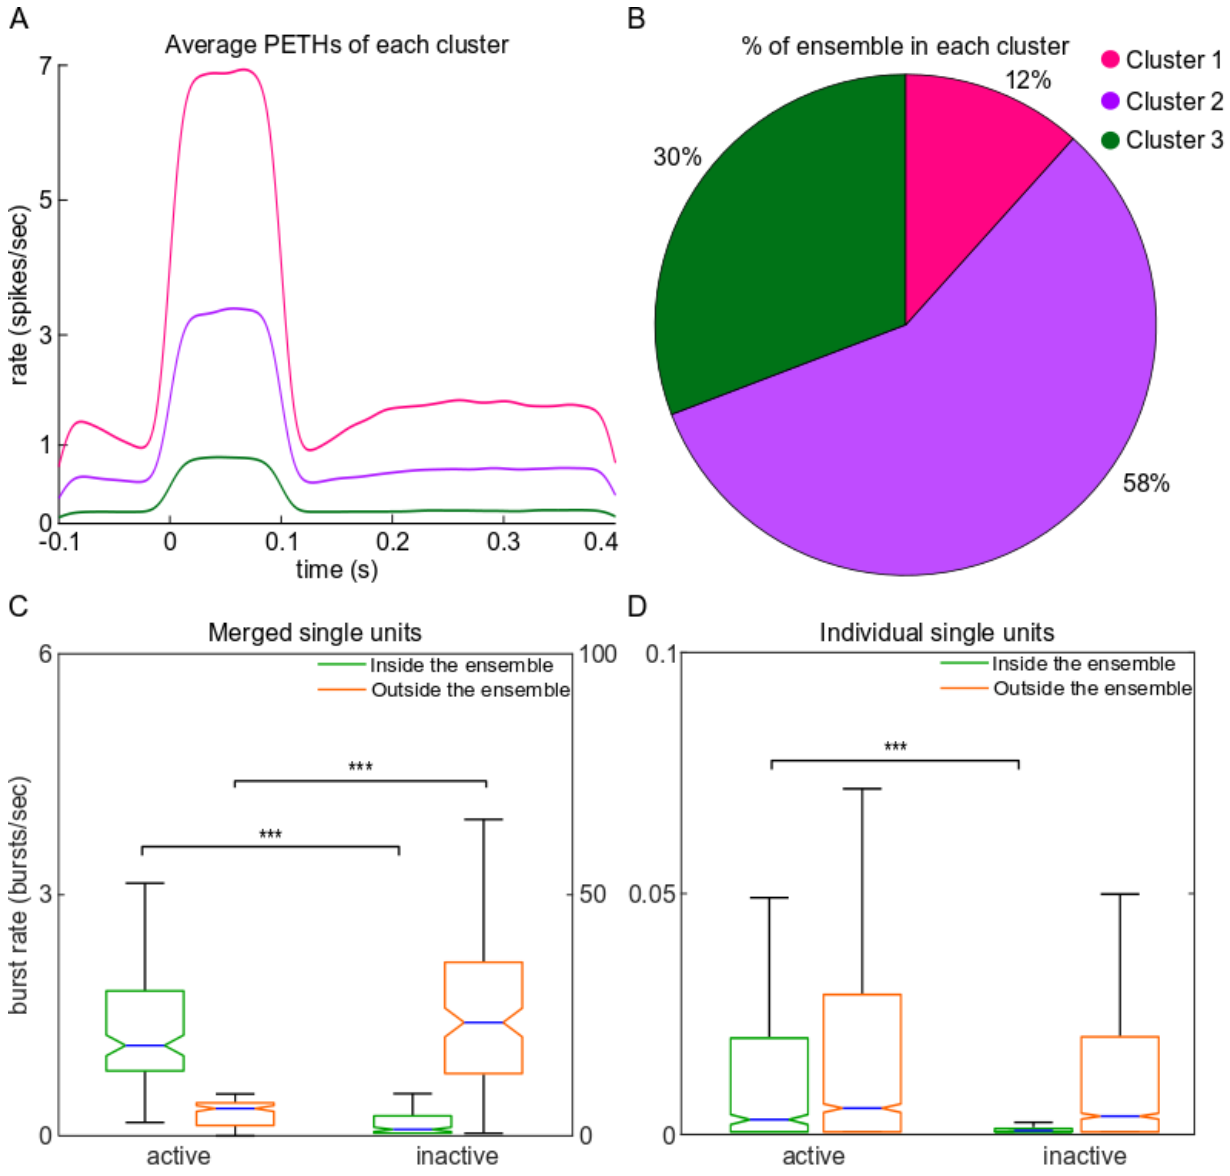

Figure 2. LC ensembles are characterized by diverse changes in spike rate and burst firing. **(A)** Average PETHs of the ensembles in the same cluster. The zero time on the x-axis is the ensemble active time. Ensembles were grouped into 3 clusters, or PETH types, that increased their firing rate to different degrees when the ensemble was active. **(B)** The pie chart illustrates the percentage of ensembles in each PETH cluster. Most ensembles had a medium or low change in single unit spike rate. **(C, D)** The box plots show the distribution of the burst rates for the single units inside the ensemble (green) and outside the ensemble (orange). The result is shown separately for the spike trains merged across single units (C) and individual single units (D). Single units inside the ensemble burst more frequently during ensemble active times in both cases. For the single units outside the ensemble there was an increased burst rate for units outside of the ensemble during periods when the ensemble was inactive, but only when single units were merged.

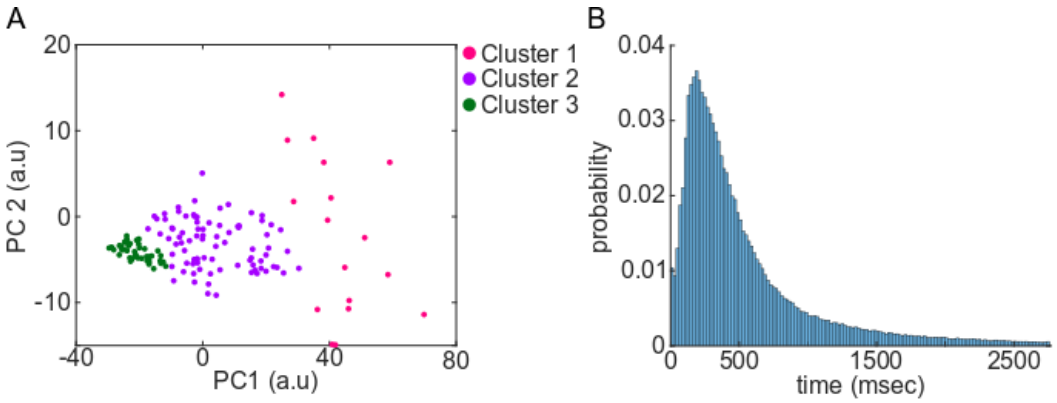

Supplementary Figure 3. Supporting data showing the clustering of PETHs and the justification of an 80 msec inter-spike interval to define bursting. (A) The scatter plot shows the projections of the PETHs into two dimensions (PC1, PC2) using PCA. The first two principle components explained more than 95% of the variance. Three non-circular masses were clustered using GMM. Data points falling into each cluster are color coded separately. **(B)** Probability distribution of the pool inter spike interval for all the recorded single units.

LC ensembles are susceptible to self-inhibition, but laterally-inhibit sparsely

Local noradrenergic inhibition is a prevalent determinant of the spiking patterns of individual LC neurons via self-inhibitory and lateral inhibition neuronal circuit motifs (34, 35). LC neurons are large (25 um soma) and densely packed with numerous close proximity dendrites (47–49), which are the site of alpha-2 self-receptors that can mediate noradrenergic self-inhibition and lateral inhibition (50, 51). Thus, the activation of multiple LC neurons and the volume transmission of noradrenaline across closely packed dendrites should inhibit a large number of neurons in the LC. Indeed, highly localized direct electrical stimulation in the LC inhibits all recorded neurons around 200 um from the stimulation site (25). This 200 um “halo” of inhibition evoked by local population activity may be considered nearly global within the LC nucleus given the small size of the LC. Therefore, current evidence suggests that LC ensemble activations – the simultaneous activation of more than one LC single unit – should be associated with a large spread of lateral inhibition across many single units.

The NMF method of detecting ensembles is well suited for assessing ensemble activation-evoked self-inhibition and lateral inhibition because the method provides the times at which ensembles are activated and inactivated. We assessed self-inhibition of LC ensembles by examining LC ensemble activation timing auto-correlograms (**Figure 3A**). We assessed lateral inhibition of LC ensembles by measuring the cross-correlations between LC ensemble activation times (**Figure 3B**). We expected LC ensembles to self-inhibit because the synchronous spiking of single units

within an ensemble should be followed by a self-inhibition of each single unit that, which is synchronized across the single units. We also expected to observe lateral inhibition between LC ensembles, as the collective spiking of neurons within an ensemble should lead to a synchronous release of norepinephrine that can inhibit other LC single units.

We found that self-inhibition occurred in the majority (62%) of the 146 ensembles. Inhibition lasted less than 300 msec for nearly all of the ensembles with significant inhibition (98%) and peaked at 100 msec (**Figure 3E**). When we considered lateral inhibition between pairs of LC ensembles, we observed 44% (out of 790 total ensemble-pairs) had at least one significant inhibitory interaction with another ensemble. The histogram showing the timing of significant lateral inhibitory interactions between ensemble-pairs has a peak at ±300 msec (**Figure 3F**). Overall, these analyses demonstrate some similarities between LC functional motifs for single units and ensembles. We show that LC ensembles tend to inhibit themselves. Moreover, we show that some ensembles laterally-inhibit other ensembles. However, in many cases, the activation of an LC ensemble does not cause a global "halo" of surrounding inhibition across the LC.

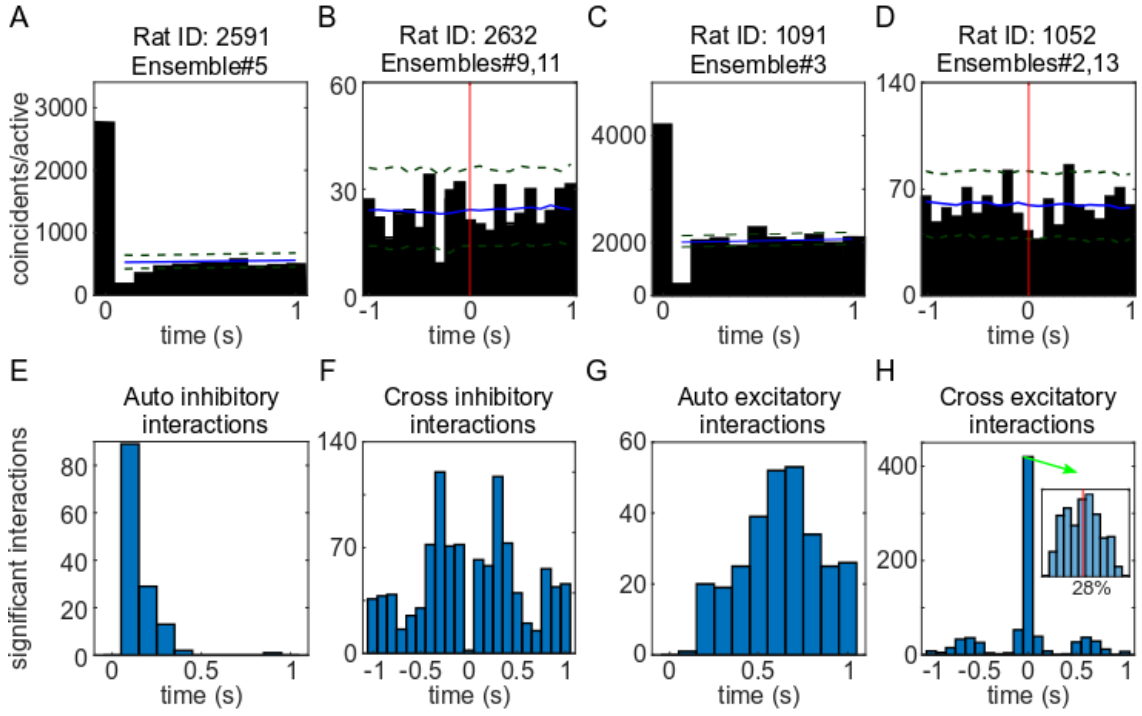

Figure 3. LC ensembles exhibited self-inhibition and lateral inhibition and patterned self-excitation and co-excitation. (A-D) Four examples of self-interactions (A, C) and cross-interactions (B, D). Significant excitatory or inhibitory interactions were defined as those that crossed the upper (excitation) or lower (inhibition) bounds of the 1% pairwise maximum or minimum threshold (dashed green lines) calculated using surrogate data from jittering ensemble

active times. The solid blue line shows the average of the surrogate correlograms. **(E)** Histogram of significant self-inhibition times for all the ensembles with significant self-inhibition. The plot shows that self-inhibition in almost all cases (98%) lasts less than 300 msec. **(F)** Histogram of significant lateral-inhibition times for all ensemble-pairs that exhibited significant lateral-inhibition. The histogram shows a peak at ±300 msec. **(G)** Histogram of significant self-excitation among ensembles. Self-excitation happens after 300 msec in 73% of the ensembles. **(H)** Histogram of significant co-excitation times reveals an initial peak at time 0 and another peak around ±600 msec. The inner plot shows the histogram of synchrony index between the ensembles that had significant co-excitation at time 0. The average synchrony index is 28%, which indicates low zero-lag synchrony among the detected ensembles.

LC ensemble pairs exhibit patterned excitations

Although the LC has no intrinsic excitatory neurotransmitters, it does receive numerous extrinsic sources of excitatory input (52–60). While LC ensembles cannot directly excite one another, it is possible that one ensemble is consistently activated after another ensemble due to an extrinsic input (or inputs), which systematically pattern the activation times of LC ensembles. First, in LC ensemble activation auto-correlograms, we observed self-excitation occurring after 300 msec in 73% of the 146 ensembles (**Figure 3G**). Among 790 ensemble-pairs, we observed 64% had excitatory interactions which peaked around time zero and again around ±600 msec (**Figure 3H**). We quantified the amount of synchrony between ensembles at time zero using a synchrony index (see methods section for the equation specifying how the index was calculated). The average synchrony between the ensembles with significant co-activation at time 0 was around 28% (**Figure 3H**, inset). The low value of this index shows that ensembles that do have a significant correlation at time zero are, on average, activated together only for a limited amount of activations (on average 28%). Therefore, the ensembles identified by NMF are indeed rarely firing in zero-lag synchrony and are indeed separate ensembles. Our analyses demonstrate that LC ensembles can respond synchronously, but only rarely. Moreover, rarely occurring cross-ensemble synchrony can happen in a pattern of at least two synchronous bursts separated by 600 msec. Overall, interactions between LC ensembles are diverse and vary from a total lack of interaction to rare bouts of synchrony.

Neuromodulation by different LC ensembles is associated with diverse changes in the cortical local field potential

Having shown that the LC has unit type-specific ensembles with a range of neuronal sizes and diverse activation dynamics over time, we turned to characterizing the neuromodulatory effects

of LC ensembles on cortical activity. Prior work has demonstrated that LC neurons, when activated, cause a change in brain state during anesthesia as well as sleep (25, 61, 62). This state change is characterized by a decrease in low frequency cortical local field potential (LFP) oscillations and an increase in high frequency cortical LFP oscillations, which has been referred to as an “up state”, “activated state”, or “prolonged up state.” LC neurons project broadly throughout the forebrain and may cause global brain-wide state changes; however, there is also evidence for some targeted specificity in the projections of single LC neurons. Therefore, depending on which LC ensemble is activated, the forebrain location of neuromodulation may vary.

We studied this possibility by examining changes in prefrontal cortex LFP power evoked by activation of specific LC ensembles. We triggered the LFP on the timings of LC ensemble activations and calculated a spectrogram. We chose a time window of 400 msec before ensemble activation until 500 msec afterwards. This window was chosen for two reasons. First, it provided a good tradeoff between temporal and spectral resolution. Second, our previous analyses of cross-correlations (**Figure 3**) and durations of activation and inactivation (**Figure 1D**) show that it is unlikely for ensembles to be co-active during this window. This window ensured that changes in the cortical LFP spectrum, triggered on activation of a specific ensemble, could not be due to another simultaneously active ensemble in the recording. After calculating ensemble activation-triggered spectra, the spectra were averaged. This yielded one activation-triggered spectrogram per ensemble (N = 89 ensembles; the smaller number of ensembles is due to some rats not having a recording of the cortical LFP). Visual inspection of the LC ensemble activation-triggered spectra revealed diverse brain states in the prefrontal cortex depending on which LC ensemble was activated.

In order to determine if there were any predominate spectral patterns, we clustered the spectrograms. The optimal number of clusters was 4 (see methods for a detailed description of how the optimal number of clusters was determined). Briefly, we first varied the number of clusters from 1 to 22 (approximately one-quarter of the data) using various distance functions and then chose the number of clusters as the elbow of the curve where error dropped below 5% (**Supplementary Figure 4A**). A sillohute analysis was used to check consistency of clustering across different distance functions (**Supplementary Figure 4B**). Together, these analyses demonstraetd that using correlation as the distance function with 4 clusters led to acceptable results. The 4 types of spectra are shown in **Figure 4A,** top row. Our results show that different sub-sets of ensembles are associated with different power spectra. The top row of **Figure 4A**

574 shows the average spectrum evoked by activation of each sub-set (or cluster) of ensembles. The
575 3 lower panels in **Figure 4A** show the activity of a single example ensemble selected from each
576 sub-set. The spike raster, average PETH, and average activation-triggered spectrogram is shown
577 for each example ensemble. The ensemble activation time (t=0 sec) is aligned across plots and
578 indiated by the red line. In contrast to the notion of population (multi-neuron) activation in the LC
579 always evoking an “activated” cortical state, the power spectra were diverse across the 4 clusters
580 of ensembles. For instance, ensembles that fell into the first type of evoked spectra (Cluster 1 in
581 **Figure 4A**, top row) were associated with the previously-reported activated state: a decrease of
582 low frequency oscillation power and increased high frequency oscilliation power. This type of
583 spectra accounted for 28% of the 89 ensembles. The second type (Cluster 2 in **Figure 4A**) of
584 spectra, which was associated with activation of 22.5% of the ensembles, was characterized by
585 an increase in middle to high frequency components of the LFP while low frequencies did not
586 change. The third type of neuromodulatory pattern (Cluster 3 in **Figure 4A**) opposed the direction
587 of the first two spectral types, in that the middle to high frequency components of the LFP were
588 decreased. This spectral pattern was associated with 22.5% of the ensembles. In all 3 of the
589 aforementioned spectral types, the change in brain state took place after the LC ensembles
590 activated. However, the last type of spectrum (Cluster 4, **Figure 4A**) was associated with a
591 change in brain state that began before the ensemble activation, which was a decrease high
592 frequency oscillations.

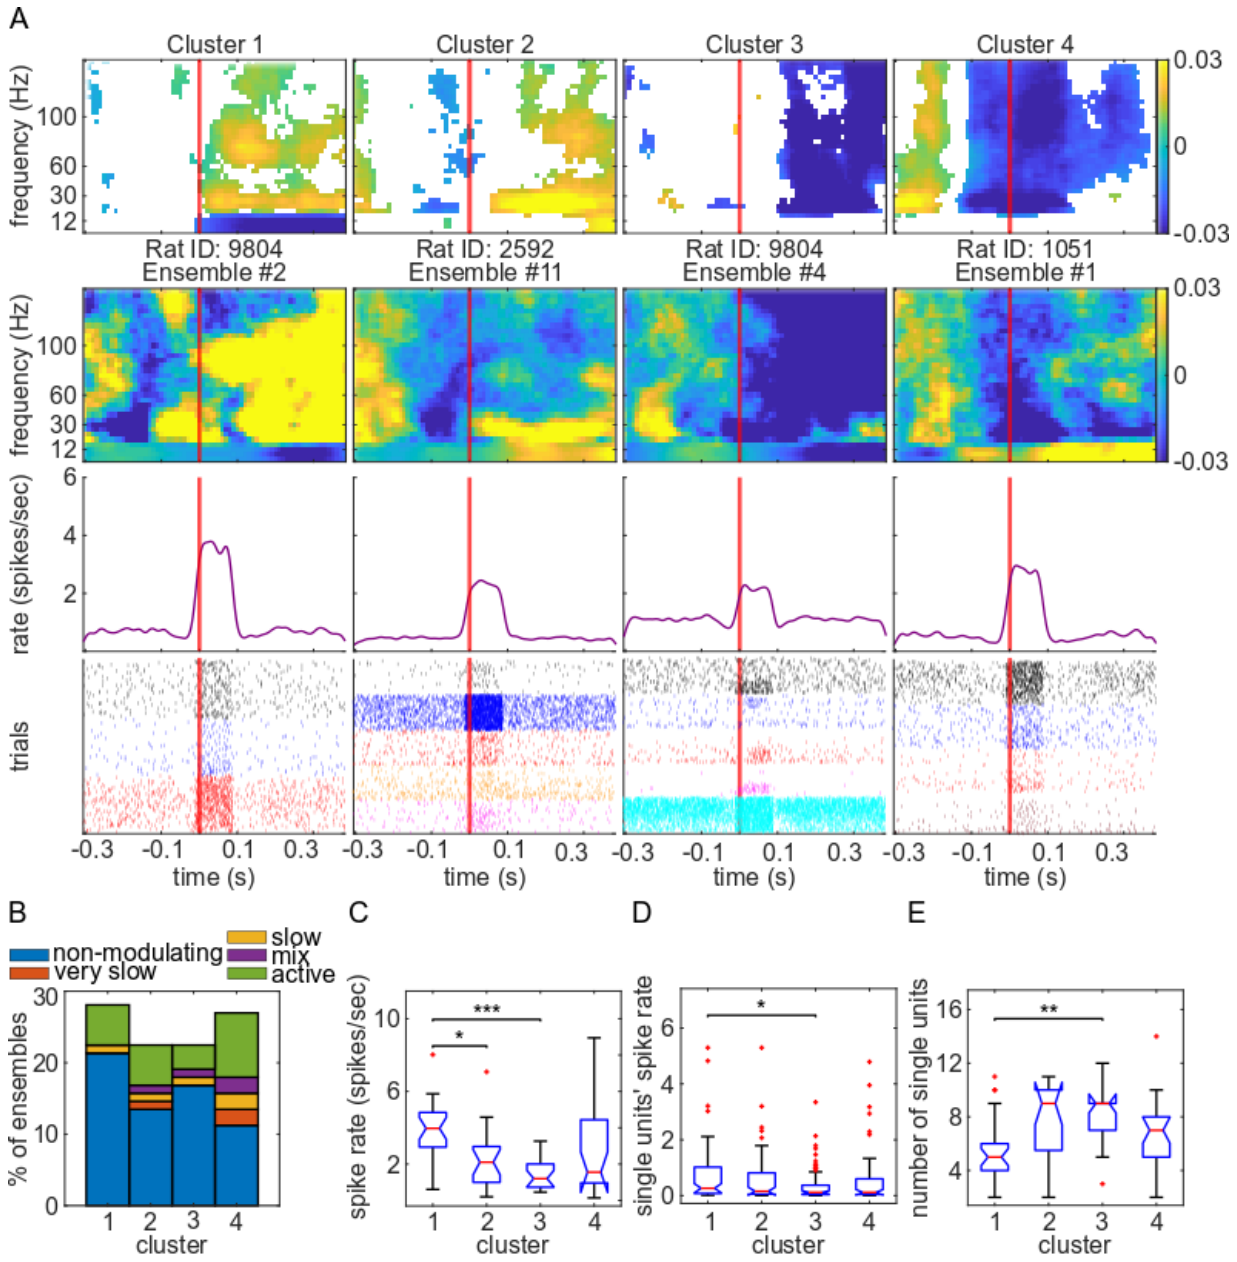

Figure 4. Activation of different LC ensembles are associated with diverse changes in prefrontal cortex LFP power spectra. (A) LFP power spectra were triggered on LC ensemble activation times. The resulting spectra were clustered into 4 spectral types, shown as 4 columns in this panel. The top row shows the average spectrogram across all ensembles associated with each type of spectrum. Only significant values are shown; other values are white. The ensemble activation time is at time 0 and is marked by a solid red line. The lower 3 rows show the activity of an example ensemble from each cluster and the ensemble activation-triggered spectrogram for that ensemble. The example spectra show significant and non-significant values. In all plots, the zero time is the ensemble activation time, which is marked by a red line. **(B)** A histogram of the percentage of the ensembles in each cluster according to the cortical state in which the ensembles were preferentially active. The blue bar shows ensembles with no preferred state). **(C)**

The box plots show the distributions of PETH maximal spike rates of the ensembles in each spectral cluster. There was a significant different in PETH firing rate between clusters 1 and 2, as well as between clusters 1 and 3. **(D)** The boxplots illustrate the distribution of the spike rate averaged across single units within the ensembles and separating the ensembles by spectral cluster. A significant difference was observed only between clusters 1 and 3. **(E)** The boxplots show the distributions of the number of single units within the ensembles for the different spectral clusters. A significance different was found only between spectral clusters of 1 and 3.

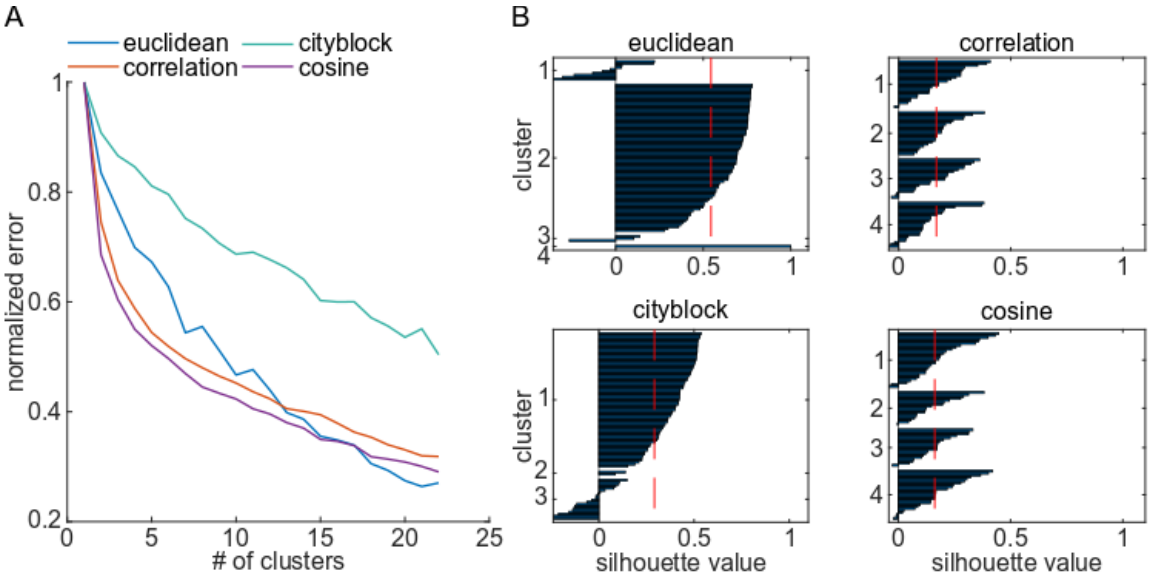

Supplementary Figure 4. The result of analyses supporting the determination of the best criteria for spectral clustering. (A) The normalized error (error divided by the maximum error) of the k-means clustering of the ensemble activation-triggered spectra versus the number of clusters. Four different distance measures were assessed and each is plotted in a different color. **(B)** Each panel shows the result of the silhouette analyses on the chosen number of clusters for four different distance measures. The optimal distance was selected based on both the uniformity in each cluster (the width of the bar plots) and the average silhouette value (the dashed red line).

One explaination for the diversity of brain states associated with activation of different LC enembles is specific ensembles are active only during particular brain states. Under urethane anesthesia, cortical LFP can be characterized by predominantly very slow (<1 Hz) or slow (<4 Hz) oscillations, or an “activated” state with reduced low frequency oscillations and increased high frequency (>20 Hz) oscillation amplitude (33). Most previous studies have shown that LC units tend to fire less in the slow wave state of the cortical activity (both in the natural sleep and anesthetized states) (12, 26, 63, 64). We checked to see whether this is the case for LC ensembles and if this might explain the different brain states associated with activation of different ensembles. We divided the cortical LFP into windows (duration, 7.5 seconds) and assigned a state identity to each window very slow, slow, mixture of slow and activated, or activated). A

Bayesian procedure was used to assess whether ensembles tend to be active in a specific state or not. Overall we found that most ensembles (63%) do not prefer being active during a specific brain state. The remaining ensembles were preferentially active in the activated state (24% of the all ensembles). A small number of ensembles were active primarily during other states (4% in the mixed state, 6% in the slow state, and 3% in the very slow state). These results were broken down by the number of ensembles in each power spectrum cluster (**Figure 4B**). Although some LC ensembles may on a large time scale prefer specific brain states for being active, their transient relationship with brain states (as shown in **Figure 4A**) is still diverse. The observed LC ensemble activity-triggered cortical state changes are likely not due to LC ensembles being preferentially active during a particular brain state.

Another potential explaination for the different brain states associated with various LC ensembles is that the size of the ensemble (i.e., number of single units within the ensemble) or the firing rate of the ensemble is related to the state shown in **Figure 4A**. We found that the median peak population firing rate differed by spectral cluster (Kruskal-Wallis test, p = 0.0003, ω^2 = 0.9633, χ^2 = 18.82) and post-hoc tests showed that only cluster 1 was different from 2 and 3 (**Figure 4C**). We also examined the firing rate of the single units in the ensembles and found a difference (Kruskal-Wallis test, p = 0.0334, ω^2 = 0.9871, χ^2 = 8.71), which was only between clusters 1 and 3 (**Figure 4D**). Finally, we examined the size of the ensemble and again found a difference (Kruskal-Wallis test, p = 0.0029, ω^2 = 0.9608, χ^2 = 13.97), that was again only between clusters 1 and 3 (**Figure 4E**). These results demonstrate that, while cluster 1 and 3 differ, there is no systematic relationship between the characteristics of an ensemble and the brain state change involving increased high frequency LFP oscillations versus increased low frequency LFP oscillations. Overall, our results demonstrate that brain states depend on which LC ensembles are active; moreover, the brain state change is not always a simple “activation” of the cortex.

Discussion

In this study, we used NMF to detect LC ensembles and resolve their temporal dynamics in relation to simultaneously recorded brain states in the urethane-anesthetized rat. Our findings lend support to the idea that the LC consists of cell type-specific ensembles (26, 31). We found that ensembles have a stable membership (i.e., units prefer one ensemble). Ensembles are rarely co-active. Analysis of the temporal dynamics of ensemble activity revealed that LC ensembles

662 are self-inhibitory and also laterally-inhibit one another. Self- and lateral-inhibition may be
663 physiological factors that promote activation of different ensembles at different times, thus
664 explaining the rare co-activation of ensembles. We found that LC ensemble activity was
665 associated with burst firing (i.e., <80 msec inter-spike interval) both a population level across units
666 in the ensemble and for individual units in the ensemble. The former may be relevant to tightly-
667 timed norepinephrine release in multiple brain regions innervated by the units in an ensemble,
668 whereas the latter may promote increased norepinephrine release from single units (46, 65).
669 Finally, we observed different LFP power spectra in the prefrontal cortex depending on which
670 ensemble was active. Overall, our results demonstrate that spontaneous activation of different LC
671 ensembles is associated with a different brain states. Identifying the neurons that control global
672 brain state has been a fundamental topic of research that has largely focused on diffusely-
673 projecting neuromodulatory centers, such as the LC (1, 2, 4). Here, we have shown that
674 spontaneous LC activity consists of independently active ensembles that are not associated with
675 a global change in brain state, that is, a similar state in prefrontal cortex regardless of which LC
676 ensemble is active.

677 Our demonstration of diverse, non-global neuromodulation by the LC, while limited to the
678 anesthetized rat and prefrontal cortex, strongly indicates that future research should attempt
679 whole-brain state recordings (e.g., high-density EEG) simultaneously with large-scale LC single
680 unit recordings, especially in the awake and naturally sleeping organism. We recorded in a single
681 cortical region (prelimbic division of medial prefrontal cortex) that is estimated to receive
682 projections from approximately 61 to 65% of LC neurons in the rat (66, 67). The LC also densely
683 innervates the thalamus, which itself projects broadly to the cortex (5). Therefore, it seems likely
684 that activation of any sub-set of LC neurons should be associated with similar brain states in the
685 prefrontal cortex. Our finding of LC ensemble-specific modulation in the prefrontal cortex
686 demonstrates that this is not the case. However, sampling from additional cortical regions with
687 high-density EEG simultaneously would provide a much-needed account of how many and which
688 cortical regions can be simultaneously modulated by a single LC ensemble. Experiments building
689 upon the present work should also examine LC ensemble-specific brain states in awake and
690 naturally sleeping organisms (e.g., rodents, non-human primates); however, it is worth noting that
691 chronic LC recordings in awake organisms are challenging due to damage at the electrode-tissue
692 interface caused by movement of the brainstem with each breath and body movement.
693 Experiments in the awake organism could reveal an LC ensemble code that is different from what
694 has been observed under anesthesia in the present study and in prior work (26). For example, it

is possible that, at points of behavioral transition (e.g., from sleep to wakefulness), all LC ensembles become transiently co-active to adjust global brain state, in contrast to our findings under anesthesia that ensemble co-activation is rare (28% of the time on average among 64% of ensemble-pairs). Synchronous ensemble activity among cholinergic neurons has been observed at behavioral transitions (locomotion onset) in mice (68). Our finding that ensemble co-activation sometimes occurs (28% of the time) suggests that LC ensembles could affect global brain state, at least transiently. Future research will need to focus on the extent to which ensemble co-activation occurs in the awake organism, especially at points of behavioral transition. It will also need to assess brain states associated with LC ensemble activity using a more widespread sampling of cortical activity (i.e., high-density EEG recordings).

Our conclusion that LC neuromodulation can be non-global requires LC ensembles to become active at distinct times. We observed that LC ensembles were active independently from one another. This configuration of activity must involve LC afferents or intra-LC neurotransmitters that enforce the independence between ensembles. It is important to characterize afferent sources and their neurochemistry, as well as the intra-LC neurotransmitters that structure LC ensemble membership and ensemble activation times because targeting these sources with pharmacological manipulation may enable the regulation of brain states related to sleep, learning and memory, and cognition. Our analyses of LC ensemble activity dynamics suggest roles for excitatory afferents, as well as intra-LC norepinephrine volume neurotransmission in maintaining the independence between LC ensembles and structuring their activation timing.

The idea the intra-LC noradrenergic self-inhibition and lateral-inhibition may enable LC ensembles to activate independently is supported by analysis of LC ensemble auto-correlograms and cross-correlograms and the previously-reported durations of intra-LC inhibition. The auto-correlation of ensemble activity frequently showed that self-inhibition was maximal 100 msec after ensemble activation and lasted only a few hundred msec. This duration is similar to spontaneous self-inhibition of LC single units (35). It is also consistent with the duration of self-inhibition after a single current pulse evokes an increase in spontaneous firing (25). The synchronized activation of all single units in an ensemble likely causes a post-activation inhibition that is also synchronized across units in an ensemble. Self-inhibition by LC ensembles may rapidly curtail activity to prevent synchrony across ensembles and promote independence between ensembles. Lateral inhibition may offer a similar constraint on LC ensemble activation due to the intra-LC volume transmission of norepinephrine released during neuronal activity, which would stimulate alpha-2 noradrenergic receptors on neurons in other ensembles. We observed lateral inhibition between LC ensembles,

which was apparent as decreased spike counts in ensemble-pair cross-correlograms. Although activation of multiple LC neurons in an ensemble might be expected to produce a "halo" of surrounding inhibition across the LC, we observed lateral inhibition among only 44% of ensemble-pairs. Overall, these analyses demonstrate some similar neurobiological motifs that apply to both single LC neurons and LC ensembles and may cause LC ensembles to activate at distinct times.

Despite our finding that LC ensembles activated independently from one another, we also observed ensemble-pair synchrony, albeit rarely (among 64% of ensemble-pairs, which co-activated only 28% of the time). Additionally, self-excitation was observed for single ensembles. The drive over LC ensemble activations must be controlled by excitatory afferents because LC neurons cannot directly excite other LC neurons since they do not contain excitatory neurotransmitters. The LC receives numerous excitatory afferents that release glutamate, orexin, corticotrophin-releasing factor, histamine, or acetylcholine – all of which have an excitatory effect on LC neurons (52–60). Future investigation of excitatory afferents will be necessary to understand how LC ensemble activations are "controlled" in a way that maintains their independent activation times, as well as their rare co-activations.

In this study, we addressed the critical question of whether different LC ensembles would be associated with the same global brain state. Prior work using direct electrical or optogenetic stimulation has suggested that LC activation produces a canonical "activated" brain state, namely, a decrease in the amplitude of low frequency (delta/theta, <10 Hz) and increased amplitude of high frequency (beta/gamma, >20 Hz) EEG/LFP oscillations (9, 16, 24, 25). Spontaneously increased LC multi-unit activity is also followed by an activated brain state (69). Such prior works have either activated the entire LC *en masse* with stimulation or they have studied the activity of many single units combined. When we use NMF to deconstruct LC population activity into discrete ensembles, we do not observe an LC-associated activated cortical state. Instead, we observed 4 types of cortical state depending on which LC ensemble was active. Two of these states resembled the previously-reported activated state, but with an important difference. While one state was characterized by a decrease in delta/theta power and an increase in beta/gamma power, the other state was characterized by a pronounced increase in beta (20 – 30 Hz) power and no change to the lower frequency bands. The other two states involved a decrease in beta/gamma power without a change in the low frequencies, which is not consistent with prior work demonstrating the activated state following LC spiking. These various brain states were not related to the size of the ensemble or its population firing rate. It is possible that differences in projection targets across LC enembles contributes to the differences in observed state (although

761 we note that ~65% of LC single units project to the cortical site recorded). One implication of LC
762 ensemble-specific control over diverse brain states is that it may allow the LC to control specific
763 brain states and behaviors. Distinct brain states (e.g., beta activation, beta/gamma activation, or
764 beta/gamma de-activation) may have different effects on how cortical neurons respond to stimuli
765 and, consequentially, on state-dependent perceptual ability, decision-making, and action speed.
766 For example, LC ensembles associated with a beta activation state may have different effects on
767 perception compared to LC ensembles associated with a beta/gamma de-activation state. These
768 features of the LC may allow it to function as a controller over specific brain states and behaviors,
769 as well as a global arousal signal.

770 Overall, our findings support the conclusion that the LC may provide targeted neuromodulation of
771 cortical state using an ensemble code. Our results question whether the LC is an undifferentiated
772 state controller that evokes a homogenous activated state across brain regions (1, 2, 9). They
773 also provide a new perspective on prior work demonstrating that norepinephrine release in
774 monkey visual cortex can be localized to specific orientation columns and dependent upon
775 presentation of a stimulus in the preferred orientation (70). In this case, the authors speculated
776 that localized noradrenergic neuromodulation might be due to gating of release from LC axons.
777 However, our results suggest another speculative explanation, which is that local norepinephrine
778 release could be controlled by different LC ensembles that project to different orientation columns.
779 Finally, our results may explain one possible route through which local adjustments to brain state,
780 such as “local sleep”, could occur (71).

781

782 **Acknowledgements**

783 This research was supported by the Helsinki Institute of Life Science (NKT) and the Department
784 of Physiology of Cognitive Processes in the Max Planck Institute for Biological Cybernetics (IZ,
785 NKL, NKT). SP was supported by the Fondazione Caritro and by a SFARI explorer grant (Grant
786 no. 602849).

787 **Author Contributions**

788 Conceptualization: NKT, SP; Methodology: NKT, SN, SP; Formal analysis: SN; Investigation: IZ,
789 NKT; Resources: NKL, SP; Writing: NKT, SN, SP; Visualization: SN; Supervision: NKT, SP;
790 Funding acquisition: NKL, NKT, SP.

Declaration of Interests

The authors declare no competing interests.

References

1. K. D. Harris, A. Thiele, Cortical state and attention. *Nat Rev Neurosci* **12**, 509-523 (2011).

2. S.-H. Lee, Y. Dan, Neuromodulation of brain states. *Neuron* **76**, 209-222 (2012).

3. M. J. McGinley, *et al.*, Waking State: Rapid Variations Modulate Neural and Behavioral Responses. *Neuron* **87**, 1143-1161 (2015).

4. M. Steriade, D. McCormick, T. Sejnowski, Thalamocortical oscillations in the sleeping and aroused brain. *Science* **262**, 679-685 (1993).

5. L. W. Swanson, B. K. Hartman, The central adrenergic system. An immunofluorescence study of the location of cell bodies and their efferent connections in the rat utilizing dopamine-beta-hydroxylase as a marker. *J Comp Neurol* **163**, 467-505 (1975).

6. J. H. Morrison, M. E. Molliver, R. Grzanna, Noradrenergic innervation of cerebral cortex: widespread effects of local cortical lesions. *Science* **205**, 313-316 (1979).

7. B. D. Waterhouse, D. J. Woodward, Interaction of norepinephrine with cerebrocortical activity evoked by stimulation of somatosensory afferent pathways in the rat. *Exp Neurol* **67**, 11-34 (1980).

8. D. A. McCormick, Neurotransmitter actions in the thalamus and cerebral cortex and their role in neuromodulation of thalamocortical activity. *Prog Neurobiol* **39**, 337-388 (1992).

9. M. Steriade, F. Amzica, A. Nuñez, Cholinergic and noradrenergic modulation of the slow (approximately 0.3 Hz) oscillation in neocortical cells. *Journal of Neurophysiology* **70**, 1385-1400 (1993).

10. B. D. Waterhouse, H. C. Moises, D. J. Woodward, Phasic activation of the locus coeruleus enhances responses of primary sensory cortical neurons to peripheral receptive field stimulation. *Brain Res* **790**, 33-44 (1998).

11. D. M. Devilbiss, B. D. Waterhouse, The Effects of Tonic Locus Ceruleus Output on Sensory-Evoked Responses of Ventral Posterior Medial Thalamic and Barrel Field Cortical Neurons in the Awake Rat. *J Neurosci* **24**, 10773-10785 (2004).

12. G. Aston-Jones, F. Bloom, Activity of norepinephrine-containing locus coeruleus neurons in behaving rats anticipates fluctuations in the sleep-waking cycle. *J Neurosci* **1**, 876-886 (1981).

13. G. Aston-Jones, J. Rajkowski, P. Kubiak, T. Alexinsky, Locus coeruleus neurons in monkey are selectively activated by attended cues in a vigilance task. *J Neurosci* **14**, 4467-4480 (1994).

14. J. Rajkowski, H. Majczynski, E. Clayton, G. Aston-Jones, Activation of monkey locus coeruleus neurons varies with difficulty and performance in a target detection task. *J Neurophysiol* **92**, 361-371 (2004).

15. J. Rajkowski, P. Kubiak, G. Aston-Jones, Locus coeruleus activity in monkey: Phasic and tonic changes are associated with altered vigilance. *Brain Res Bull* **35**, 607-616 (1994).

16. M. E. Carter, *et al.*, Tuning arousal with optogenetic modulation of locus coeruleus neurons. *Nat Neurosci* **13**, 1526-1533 (2010).

17. C. M. Constantinople, R. M. Bruno, Effects and mechanisms of wakefulness on local cortical networks. *Neuron* **69**, 1061-1068 (2011).

18. R. L. Navarra, B. D. Clark, G. A. Zitnik, B. D. Waterhouse, Methylphenidate and atomoxetine enhance sensory-evoked neuronal activity in the visual thalamus of male rats. *Exp Clin Psychopharm* **21**, 363-374 (2013).

19. P.-O. Polack, J. Friedman, P. Golshani, Cellular mechanisms of brain state-dependent gain modulation in visual cortex. *Nat Neurosci* **16**, 1331-1339 (2013).

20. A. R. O. Martins, R. C. Froemke, Coordinated forms of noradrenergic plasticity in the locus coeruleus and primary auditory cortex. *Nat Neurosci* **18**, 1483-1492 (2015).

21. N. K. Totah, N. K. Logothetis, O. Eschenko, Atomoxetine accelerates attentional set shifting without affecting learning rate in the rat. *Psychopharmacology* **232**, 3697-3707 (2015).

22. M. Lovett-Barron, *et al.*, Ancestral Circuits for the Coordinated Modulation of Brain State. *Cell* **171**, 1411-1423.e17 (2017).

23. H. Gelbard-Sagiv, E. Magidov, H. Sharon, T. Hendler, Y. Nir, Noradrenaline Modulates Visual Perception and Late Visually Evoked Activity. *Curr Biol* **28**, 2239-2249 (2018).

24. H. Hayat, *et al.*, Locus-coeruleus norepinephrine activity gates sensory-evoked awakenings from sleep. *Biorxiv* (2019).

25. A. Marzo, N. K. Totah, R. M. Neves, N. K. Logothetis, O. Eschenko, Unilateral electrical stimulation of rat locus coeruleus elicits bilateral response of norepinephrine neurons and sustained activation of medial prefrontal cortex. *J Neurophysiol* **111**, 2570-2588 (2014).

26. N. K. Totah, R. M. Neves, S. Panzeri, N. K. Logothetis, O. Eschenko, The Locus Coeruleus Is a Complex and Differentiated Neuromodulatory System. *Neuron* **99**, 1055-1068.e6 (2018).

27. S. Nakamura, Some electrophysiological properties of neurones of rat locus coeruleus. *J Physiology* **267**, 641-658 (1977).

28. T. Sakaguchi, S. Nakamura, The mode of projections of single locus coeruleus neurons to the cerebral cortex in rats. *Neuroscience* **20**, 221-230 (1987).

29. P. W. Howorth, *et al.*, Retrograde Viral Vector-Mediated Inhibition of Pontospinal Noradrenergic Neurons Causes Hyperalgesia in Rats. *J Neurosci* **29**, 12855-12864 (2009).

30. J. M. Kebschull, *et al.*, High-Throughput Mapping of Single-Neuron Projections by Sequencing of Barcoded RNA. *Neuron* **91**, 975-987 (2016).

31. N. K. B. Totah, N. K. Logothetis, O. Eschenko, Noradrenergic ensemble-based modulation of cognition over multiple timescales. *Brain Res* **1709**, 50-66 (2018).

32. D. D. Lee, H. S. Seung, Learning the parts of objects by non-negative matrix factorization. *Nature* **401**, 788-791 (1999).

865 33. E. A. Clement, *et al.*, Cyclic and Sleep-Like Spontaneous Alternations of Brain State Under
866 Urethane Anaesthesia. *Plos One* **3**, e2004 (2008).

867 34. G. K. Aghajanian, J. M. Cedarbaum, R. Y. Wang, Evidence for norepinephrine-mediated
868 collateral inhibition of locus coeruleus neurons. *Brain Res* **136**, 570-577 (1977).

869 35. M. Ennis, G. Aston-Jones, Evidence for self- and neighbor-mediated postactivation inhibition
870 of locus coeruleus neurons. *Brain Res* **374**, 299-305 (1986).

871 36. A. Onken, *et al.*, Using Matrix and Tensor Factorizations for the Single-Trial Analysis of
872 Population Spike Trains. *Plos Comput Biol* **12**, e1005189 (2016).

873 37. W. M. Rand, Objective Criteria for the Evaluation of Clustering Methods. *J Am Stat Assoc* **66**,
874 846 (1971).

875 38. S. Fujisawa, A. Amarasingham, M. T. Harrison, G. Buzsáki, Behavior-dependent short-term
876 assembly dynamics in the medial prefrontal cortex. *Nat Neurosci* **11**, 823-833 (2008).

877 39. H. Hotelling, Analysis of a complex of statistical variables into principal components. *J Educ*
878 *Psychol* **24**, 417-441 (1933).

879 40. G. McLachlan, D. Peel, Wiley Series in Probability and Statistics (2000)
880 https:/doi.org/10.1002/0471721182.

881 41. P. Mitra, H. Bokil, *Observed Brain Dynamics* (2007)
882 https:/doi.org/10.1093/acprof:oso/9780195178081.001.0001.

883 42. D. Arthur, S. Vassilvitskii, k-means++: The Advantages of Careful Seeding. 1027-1035 (2007).

884 43. P. J. Rousseeuw, Silhouettes: A graphical aid to the interpretation and validation of cluster
885 analysis. *J Comput Appl Math* **20**, 53-65 (1987).

886 44. Y. Benjamini, Y. Hochberg, Controlling the False Discovery Rate: A Practical and Powerful
887 Approach to Multiple Testing. *J Royal Statistical Soc Ser B Methodol* **57**, 289-300 (1995).

888 45. S. Holm, A simple sequentially rejective multiple test procedure. **6**, 65-70 (1979).

889 46. G. Aston-Jones, M. Segal, F. E. Bloom, Brain aminergic axons exhibit marked variability in
890 conduction velocity. *Brain Res* **195**, 215-222 (1980).

891 47. R. Grzanna, M. E. Molliver, The locus coeruleus in the rat: An immunohistochemical
892 delineation. *Neuroscience* **5**, 21-40 (1980).

893 48. N. Shimizu, S. Ohnishi, K. Satoh, M. Tohyama, Cellular Organization of Locus Coeruleus in
894 the Rat as Studied by Golgi Method. *Arch Histol Japon* **41**, 103-112 (1978).

895 49. L. W. Swanson, The locus coeruleus: a cytoarchitectonic, Golgi and immunohistochemical
896 study in the albino rat. *Brain Research* **110**, 39-56 (1976).

897 50. A. Lee, D. L. Rosin, E. J. V. Bockstaele, α2A-adrenergic receptors in the rat nucleus locus
898 coeruleus: subcellular localization in catecholaminergic dendrites, astrocytes, and presynaptic
899 axon terminals. *Brain Res* **795**, 157-169 (1998).

900 51. H.-P. Huang, *et al.*, Long latency of evoked quantal transmitter release from somata of locus
901 coeruleus neurons in rat pontine slices. *Proc National Acad Sci* **104**, 1401-1406 (2007).

902 52. P. G. Guyenet, G. K. Aghajanian, ACh, substance P and met-enkephalin in the locus
903 coeruleus: Pharmacological evidence for independent sites of action. *Eur J Pharmacol* **53**, 319-
904 328 (1979).

905 53. G. Engberg, T. H. Svensson, Pharmacological analysis of a cholinergic receptor mediated
906 regulation of brain norepinephrine neurons. *Journal of neural transmission* **49**, 137-150 (1980).

907 54. A. F. T. Arnsten, P. S. Goldman-Rakic, Selective prefrontal cortical projections to the region
908 of the locus coeruleus and raphe nuclei in the rhesus monkey. *Brain Res* **306**, 9-18 (1984).

909 55. G. Aston-Jones, *et al.*, Afferent regulation of locus coeruleus neurons: anatomy, physiology
910 and pharmacology. *Progress in brain research* **88**, 47-75 (1991).

911 56. R. J. Valentino, M. E. Page, A. L. Curtis, Activation of noradrenergic locus coeruleus neurons
912 by hemodynamic stress is due to local release of corticotropin-releasing factor. *Brain Res* **555**,
913 25-34 (1991).

914 57. T. L. Horvath, *et al.*, Hypocretin (orexin) activation and synaptic innervation of the locus
915 coeruleus noradrenergic system. *J Comp Neurol* **415**, 145-159 (1999).

916 58. T. M. Korotkova, O. A. Sergeeva, A. A. Ponomarenko, H. L. Haas, Histamine excites
917 noradrenergic neurons in locus coeruleus in rats. *Neuropharmacology* **49**, 129-134 (2005).

918 59. J. G. McCall, *et al.*, CRH Engagement of the Locus Coeruleus Noradrenergic System
919 Mediates Stress-Induced Anxiety. *Neuron* **87**, 605-620 (2015).

920 60. V. Breton-Provencher, M. Sur, Active control of arousal by a locus coeruleus GABAergic
921 circuit. *Nat Neurosci* **28**, 218-228 (2019).

922 61. O. Eschenko, C. Magri, S. Panzeri, S. J. Sara, Noradrenergic Neurons of the Locus Coeruleus
923 Are Phase Locked to Cortical Up-Down States during Sleep. *Cereb Cortex* **22**, 426-435 (2012).

924 62. H. Safaai, R. Neves, O. Eschenko, N. K. Logothetis, S. Panzeri, Modeling the effect of locus
925 coeruleus firing on cortical state dynamics and single-trial sensory processing. *Proc National Acad*
926 *Sci* **112**, 12834-12839 (2015).

927 63. O. Eschenko, S. J. Sara, Learning-Dependent, Transient Increase of Activity in Noradrenergic
928 Neurons of Locus Coeruleus during Slow Wave Sleep in the Rat: Brain Stem–Cortex Interplay for
929 Memory Consolidation? *Cereb Cortex* **18**, 2596-2603 (2008).

930 64. K. M. Swift, *et al.*, Abnormal Locus Coeruleus Sleep Activity Alters Sleep Signatures of
931 Memory Consolidation and Impairs Place Cell Stability and Spatial Memory. *Curr Biol* **28**, 3599-
932 3609.e4 (2018).

933 65. S. M. Florin-Lechner, J. P. Druhan, G. Aston-Jones, R. J. Valentino, Enhanced norepinephrine
934 release in prefrontal cortex with burst stimulation of the locus coeruleus. *Brain Res* **742**, 89-97
935 (1996).

936 66. D. J. Chandler, W.-J. Gao, B. D. Waterhouse, Heterogeneous organization of the locus
937 coeruleus projections to prefrontal and motor cortices. *Proc National Acad Sci* **111**, 6816-6821
938 (2014).

939 67. D. J. Chandler, C. S. Lamperski, B. D. Waterhouse, Identification and distribution of
940 projections from monoaminergic and cholinergic nuclei to functionally differentiated subregions of
941 prefrontal cortex. *Brain Res* **1522**, 38-58 (2013).

942 68. M. Howe, *et al.*, Coordination of rapid cholinergic and dopaminergic signaling in striatum
943 during spontaneous movement. *Elife* **8**, e44903 (2019).

944 69. R. M. Neves, S. van Keulen, M. Yang, N. K. Logothetis, O. Eschenko, Locus Coeruleus phasic
945 discharge is essential for stimulus-induced gamma oscillations in the prefrontal cortex. *J*
946 *Neurophysiol* **119**, 904-920 (2017).

947 70. R. Marrocco, R. Lane, J. McClurkin, C. Blaha, M. Alkire, Release of cortical catecholamines
948 by visual stimulation requires activity in thalamocortical afferents of monkey and cat. *J Neurosci*
949 **7**, 2756-2767 (1987).

950 71. V. V. Vyazovskiy, *et al.*, Local sleep in awake rats. *Nature* **472**, 443-7 (2011).

951

952