Distinct ensembles in the noradrenergic locus coeruleus evoke diverse cortical states

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

The noradrenergic locus coeruleus (LC) is crucial for controlling brain and behavioral states. While synchronous stimulation of LC neurons evokes a single activated cortical state with increased high-frequency power, little is known about how spontaneous patterns of LC population activity drive cortical states. Since LC neurons selectively project to specific forebrain regions, we hypothesized that individual LC ensembles produce different cortical states. We recorded up to 34 single units simultaneously in the rat LC and used non-negative matrix factorization to identify spontaneously activated ensembles of co-active LC neurons. The ensembles were active mostly at different times and were simultaneously active only rarely. We assessed cortical state in area 24a by examining local field potential power spectrograms triggered on activations of individual LC ensembles. We observed four spectrotemporally-distinct cortical states associated with activation of specific LC ensembles. Thus, distinct spontaneously active LC ensembles contribute to unexpectedly diverse cortical states.

Introduction

Flexible behavior is associated with transitions across diverse cortical states. For example, various states of wakefulness, perceptual ability, and behavioral activity are associated with different cortical states each with its own clear pattern of neural oscillations and synchronization properties (Harris and Thiele, 2011; McGinley et al., 2015; McCormick et al., 2020). Moreover, behavioral state transitions, such as waking from sleep or entering a state of heightened stress and reacting more quickly to stimuli, are associated with cortical state transition. These changes are not necessarily driven by external stimuli. Instead, brain state can be controlled by factors internal to the organism (e.g., sleep need, perceived stress) and therefore arise from self-organized neuronal interactions. It remains unclear exactly which interactions among neurons control specific brain states.

Maintenance of brain state and transitions between states are mediated, at least in part, by the noradrenergic brainstem nucleus, locus coeruleus (LC). The LC projects globally throughout the central nervous system and releases norepinephrine to modulate neuronal excitability (Swanson and Hartman, 1975; Waterhouse and Woodward, 1980; McCormick, 1992; Devilbiss and Waterhouse, 2004). Activating neurons in the LC synchronously by direct electrical (or optogenetic) stimulation in anesthetized or sleeping animals evokes a so-called “activated” brain state, with larger oscillation power at higher frequencies and reduced slow-wave power in the mean extracellular field potential, resembling that which occurs spontaneously during the emergence from anesthesia or sleep into wakefulness (Steriade et al., 1993; Carter et al., 2010; Marzo et al., 2014; Hayat et al., 2019). Moreover, increasing noradrenaline neurotransmission can lead to those behavioral transitions that are often associated with a brain state change, such as: awakening from sleep or anesthesia, altering locomotion patterns (increased generalized movements and decreased reaction times), and improving perceptual sensitivity and attentional focus (Aston-Jones and Bloom, 1981a; Aston-Jones et al., 1994; Rajkowski et al., 1994, 2004; Carter et al., 2010; Constantinople and Bruno, 2011; Navarra et al., 2013; Polack et al., 2013; Martins and Froemke, 2015; Totah et al., 2015; Lovett-Barron et al., 2017; Gelbard-Sagiv et al., 2018; Hayat et al., 2020).

The LC has been traditionally thought to produce the activated state in the cortex via *en masse* and highly-synchronous collective firing of LC neurons (Aston-Jones and Bloom, 1981a, 1981b; Finlayson and Marshall, 1988; Ishimatsu and Williams, 1996; Alvarez et al., 2002; Chen and Sara, 2007). However, recent findings suggest that this influential view might be incomplete. We recently studied the time-averaged cross-correlation properties of over 3,000 LC single unit pairs and found a surprisingly small percentage of correlated pairs. Graph-theoretical analyses of time-average cross-correlograms suggested that the small number of correlated pairs are seemingly organized into sparse coactive ensembles. This prior work demonstrated that LC neurons clearly do not collectively fire *en masse* (Totah et al., 2018a). Furthermore, these data suggest the possibility that LC firing is organized into patterns of small ensembles of simultaneously active neurons that change from moment to moment. Given the neurochemical diversity of LC neurons, as well as diversity in their forebrain projection patterns (for review, see (Totah et al., 2018b; Chandler et al., 2019), this hypothesis raises the intriguing possibility that individual LC ensembles could have different effects on the self-organization of neuronal circuits that produce various brain states. As a result, various LC ensembles could potentially evoke distinct brain states beyond the activated state.

Here, our objectives were, first, to test the hypothesis that LC population activity consists of multiple, discrete LC ensembles each with its own evolution of activity over time and, second, to examine the relationships between brain state and LC ensemble dynamics. We addressed these open questions using a mathematical methodology called non-negative matrix factorization (NMF), which allowed us to decompose the spiking of simultaneously recorded LC single units into individual patterns of coactive neurons at any given time. In line with the predictions of our prior work (Totah et al., 2018a), we found that LC activity at any given time is formed by a small number of simultaneously active cell type-specific ensembles. Using this new approach to access the moment-to-moment changes in ensemble activity, we were able to reveal that self-inhibition and lateral-inhibition (which are common among individual LC single units (Aghajanian et al., 1977; Ennis and Aston-Jones, 1986)) also occur between LC ensembles. Although different ensembles activated primarily at different times, with evidence of some periodic structure of time-delayed excitation and inhibition of different ensembles, we observed simultaneous activation of multiple LC ensembles only rarely. We studied what – if any – brain state diversity is associated

with activation of different LC ensembles. In contrast to the canonical view that LC activation evokes a unitary activated brain state, we observed heterogenous brain states with different spectral and temporal properties that depended on which LC ensemble was active. However, when different LC ensembles were spontaneously coactivated, the associated brain states were more homogenous, in line with the prototypical activated state resulting from whole-LC stimulation. In sum, we report moment-to-moment changes in LC ensemble activity and show that spontaneous activation of separate sets of ensembles are associated with diverse cortical states.

Results

In order to study the temporal dynamics of LC population activity, we recorded from many LC single units simultaneously (5 to 34 units and, on average, 19 units recorded from 15 male rats) using a silicon probe with 32 electrodes confined to the core of the LC nucleus. Probe location was verified histologically in coronal tissue sections. Neuronal identity was confirmed at the end of the experiments using intra-peritoneal injection of the alpha-2 agonist, clonidine, which inhibited spiking on all electrodes. Spikes recorded from outside the LC core would not have been inhibited due to the lack of alpha-2 adrenergic receptors in nearby brain structures (McCune et al., 1993). In order to assess brain state, a single tungsten electrode was placed in cortical area 24a (anterior cingulate cortex) (Paxinos and Watson, 2017) and the mean extracellular field potential (8 kHz lowpass filtered) was recorded in 9 of the 15 rats. Neuronal recordings were made under urethane anesthesia, a widely-used model for studying brain state transitions evoked by LC stimulation (Marzo et al., 2014; Neves et al., 2018). To date, recordings of many LC single units simultaneously in any awake organism with multi-electrode probes has been an intractable problem due to brainstem movement associated with body movement, thus necessitating the use of anesthesia to investigate the relationship between LC ensemble activity and brain state.

Ensemble detection using non-negative matrix factorization

We began by assessing if LC population activity consists of simultaneously coactive collections of single units (ensembles) and how those population activity patterns may spontaneously change from moment to moment. We detected LC ensembles with non-negative matrix factorization (NMF) on the spike counts of all simultaneously recorded single units for each animal. For this

114 analysis, we binned activity in sliding windows that were 100 msec long, which is the time scale
115 capturing the majority of the synchrony among LC single unit pairs (Totah et al., 2018a). **Figure**
116 **1A** sketches how NMF works on hypothetical single unit spiking data. NMF decomposes the
117 population matrix containing the spike counts of each single unit in each time bin as a sum of K
118 non-negative spatial modules, each multiplied by a non-negative activation coefficient. A spatial
119 module may be thought of as a specific, often-recurring, population firing pattern. Formally, it is
120 a vector specifying the relative strength of firing of each neuron within the population pattern
121 (Onken et al., 2016; Williams et al., 2018). The activation coefficient of each module at any given
122 time describes how strongly the specific population firing pattern (the module) is recruited during
123 that time bin.

124 The number of spatial modules in the data, K, is a free parameter whose choice must be informed
125 empirically. Following established procedures (Onken et al., 2016; Williams et al., 2018), we
126 determined K for each rat, by choosing a value based on two criteria. First, the chosen K explained
127 a high amount of variance in the data with few modules (i.e., the selected value of K was in the
128 "elbow" region of the reconstruction error plotted as a function of the possible number of modules,
129 which means that using higher K would have given diminishing returns in terms of data
130 reconstruction accuracy). Second, that value for K yielded a stable recovery of the spatial modules
131 from the data regardless of the random initialization of the decomposition optimization procedure
132 (see Methods for additional details and **Supplementary Figure 1**).

133 By thresholding each module to individuate which single units were significantly active within it,
134 we could separate and define the ensemble of single units that are coactive within each module.
135 An "ensemble" was thus defined as the set of single units that crossed the activation coefficient
136 threshold for a specific spatial module. This way, we associated one and only one ensemble of
137 coactive neurons to each of the K spatial modules. For simplicity, hereafter, we will thus refer to
138 NMF decomposition into and detection of ensembles as a shorthand to indicate detection of firing
139 patterns and ensembles of coactive single units.

140 By thresholding the time course of activation coefficients to distinguish the times of significant
141 recruitment of each spatial module, we defined the times of activation of each spatial module and,

thus, the activation times of each ensemble of single units (see Methods). For brevity, the activation times of spatial modules will be referred to as "ensemble activation times."

We found, across 15 rats, a total of 146 ensembles from 285 single units. Note that a single unit can potentially be active in more than one ensemble determined by NMF. **Figure 1B** shows two exemplar LC ensembles. The left panels depict the spike rasters of single units which belong to the ensemble and the right panels show spike rasters of single units outside the ensemble. Spike rasters are aligned to the ensemble activation times. These examples clearly show that single units assigned to the same ensemble increased their firing rate during ensemble activation. On the other hand, units not assigned to the ensemble maintain their ongoing pattern of activity without systematic variations. **Figure 1C** shows another example in which LC population activity was decomposed into K = 5 different ensembles. The left panel shows population activity (spiking of all simultaneously recorded single units combined) over an exemplar 2 second recording epoch. The right panel shows this population activity decomposed into the activity of the 5 coactive ensembles. The ensembles were active in most cases at different times, but in some cases (such as at time t = 0.5 s) more than one ensemble was active. This is apparent as the time of highest population activity in the left panel. Reconstructing the total population firing rate as function of time through the NMF decomposition (which essentially involves summing up the activation time courses across the 5 ensembles) returned a good approximation of the pooled population spike rate. This example is useful both to illustrate that the NMF decomposition captured the LC population firing well and that the total firing of LC populations cannot be conceptualized as a result of *en masse* firing but rather as a nuanced sequence of different ensembles activating at largely different times.

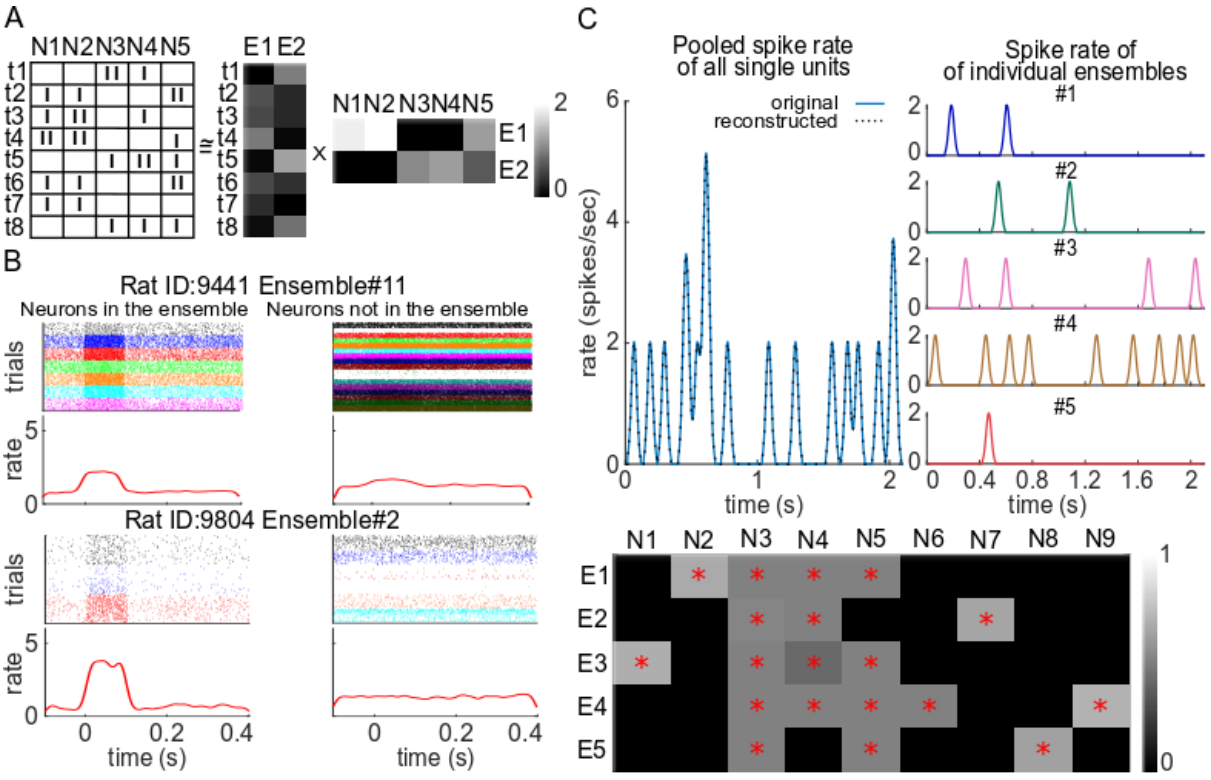

Figure 1. Ensemble detection using NMF decomposes LC population spiking into ensembles that have heterogeneous spatio-temporal properties. (A) An example of how NMF works on hypothetical data from 5 single units (N1 to N5, columns) whose spikes were binned into 8 time bins (t1 to t8, rows). The matrix of population spiking is decomposed as a sum of K non-negative spatial modules (in this example there are K= 2 spatial modules, which are plotted on the right and labeled rows E1 and E2, each representing a specific firing pattern recurring in the data and captured by the decomposition) multiplied by a non-negative activation coefficient (shown in the middle) representing the strength of recruitment of each of the two specific population spiking patterns ("modules") over time. In this example, the population spike counts (left panel) shows that units N1, N2, and N5 tend to fire concurrently with N1 and N2 firing more strongly than N5 (see time bins t2, t3, t4, t6, t7). Thus, the NMF finds a spatial module (indicated as E1) with neurons N1 and N2 firing strongly and N5 also somewhat active, but N3 and N4 not active at all. Thresholding the firing rate values of each module identifies the ensemble of units active within a module. In the case of module E1, such thresholding of the plot on the right shows that the ensemble is made of neurons N1, N2, and N5. Inspection of the plot in the middle shows that the activation coefficient of spatial module E1 is higher during those time bins (t2, t3, t4, t6, and t7) in which N1, N2 and N5 were more active (see spike times in left plot). Thresholding the activation coefficients in the middle plot detects when that module (i.e., a specific population spiking pattern) is occurring. **(B)** The spike rasters and peri-event time histograms (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 times (t = 0 sec). In these spike rasters, each ensemble activation event is a "trial." The PETHs of trial-averaged spike rate across all units

in the ensemble are shown below the rasters. The right panels depict the ensemble activation-triggered spiking of single units that were not assigned to that ensemble. 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 in any systematic way. **(C)** An example of 2 seconds of activity in a rat in which the NMF found K=5 ensembles among 9 single units (N1-9). The upper right panel shows the time course of the activation coefficients of each ensemble. The upper left panel shows that summing the ensemble activations (dotted line) reconstructs well the true LC pooled population spiking (solid blue line). The bottom panel shows the activation coefficients of each spatial module. Single units that were significantly active in a spatial module (i.e., crossed threshold) and thus formed an ensemble (E1-5) are marked by a red asterisks.

The spatio-temporal scale of LC ensemble activation is heterogeneous

Our first objective was to test the hypothesis that LC population activity consists of many discrete LC ensembles each with its own evolution of activity over time. We began by assessing the durations over which different ensembles were either spontaneously active or inactive. **Figure 2A** presents evidence that most ensembles were only transiently active for the 100 msec time bin that we used to decompose the data. However, the duration of the inactive periods varied across ensembles, such that ensembles are quiet for a wide variety of durations before being briefly active for approximately 100 msec. These findings suggest that the activation time courses vary across LC ensembles such that different ensembles are likely independently active at different times.

In order to characterize the physiological properties of LC ensembles, we next examined ensemble size and whether ensembles were spatially discrete in the LC. In **Figure 2B**, we report the relative number of single units in each ensemble relative to the total number of simultaneously recorded single units. 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 varied in size (range: 6% - 62% of the simultaneously recorded single units, see **Supplementary Figure 2A**). Single units in an ensemble were spread throughout the LC with no topographical organization (**Supplementary Figure 2B**). We assessed single unit location as the electrode which recorded the largest average spike waveform. Among 146 ensembles, only 23 ensembles had a median distance between single unit pairs of less than 50 um. The diffuse spatial arrangement of single units within ensembles detected with NMF agrees with the predictions of prior work on LC

ensembles that used graph theoretic analysis of time-averaged pairwise spike count correlations to investigate the spatial structure of synchrony within the LC (Totah et al., 2018a).

Given that NMF can identify neurons that fire in more than one ensemble (Onken et al., 2016), we also investigated how many of the single units were assigned to one ensemble, multiple ensembles, or no ensemble. Out of 285 single units, 115 single units fired as part of multiple ensembles (40.4%), 149 were active in only a single ensemble (52.3%), and the remaining 21 units did not participate in any ensemble with the other single units (**Figure 2C**). Although single units could and did take part in multiple ensembles, the probability that a neuron took part in only one ensemble was higher than the probability that a neuron took part on more than one ensemble (binomial test, $p = 0.04$).

LC ensembles were unit type-specific

Recent work has shown that two LC single unit types, termed "narrow" or "wide" type units, are distinguishable by their extracellular waveform shape (Totah et al., 2018a). We next examined whether LC single units of the same type tended to spike in the same ensemble. The recent work that first identified these LC unit types showed that units of the same type tended to form ensembles detected using graph theoretic analysis of time-averaged pairwise spike count correlations (Totah et al., 2018a). **Figure 2D** 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 ensembles are made, entirely or mostly, of units of the same type. We assessed (by random resampling) if these proportions were statistically different from what would be expected if ensembles were formed by units taken randomly regardless of their type (see Methods). Our results show that, for all rats, the hypothesis that ensembles are formed by combining units regardless of their type should be rejected ($p<0.05$), thus indicating that ensembles do not combine randomly units of different types but are instead preferentially made by units of the same type.

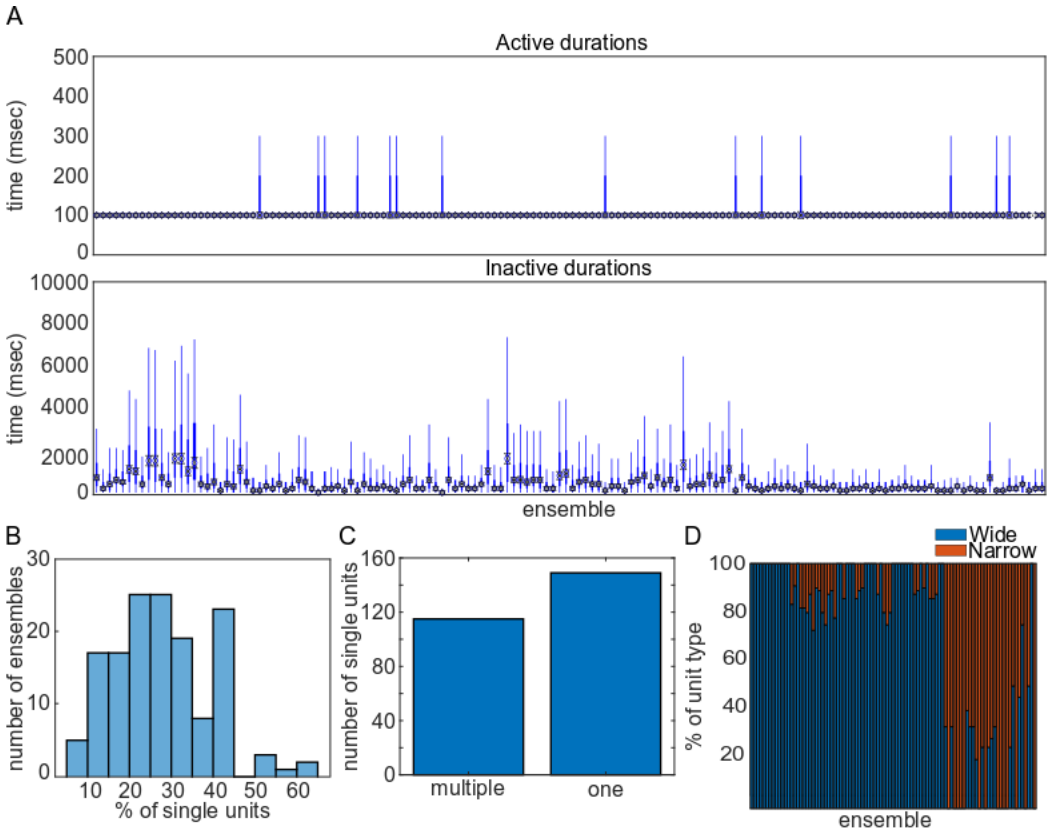

Figure 2. LC ensembles are spatio-temporally sparse and cell type-specific. (A) The box plots show the distributions of how long ensembles were consecutively active (top panel) and inactive (bottom panel). Each boxplot illustrates the distribution for one ensemble. Ensembles tend to be active for only 100 msec, but can be inactive for a wide variety of durations which yields heterogenous activation of different ensembles at different times. **(B)** 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. **(C)** 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. **(D)** The percent of each unit type (wide or narrow spike waveform) making up each ensemble is plotted across ensembles (x-axis). Ensembles are either only one type of single unit or consist of mostly a single type.

LC ensemble activation is associated with burst firing

Given our motivation to assess the association between cortical state and activation of discrete LC ensembles, it is important to determine if some ensembles fire more strongly than others or exhibit a tendency to fire in bursts, as these spiking properties might systematically vary with LC ensemble activation-associated cortical state changes.

We first characterized the firing strength of various LC ensembles. We calculated the average spike rate of all single units within the ensemble (when the ensemble was active) using peri-event time histograms (PETHs). Each event was an ensemble activation time. The PETHs were calculated 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 Principal Component Analysis (PCA) and Gaussian Mixture Model (GMM) (see Methods). When 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. These were associated with low, medium, and high changes in spike rate, but had similar activation durations (**Figure 3A**). Most ensembles (88%, green and purple in **Figures 3A** and **3B**) were characterized by a low or medium change in single unit spike rate corresponding to an increase of 1 to 3 spikes per sec (**Figure 3A**). In the maximal case, average spike rate increased by 7 spikes per sec (**Figure 3A**, magenta line), but this was the smallest group of ensembles (**Figure 3B**, magenta). Single unit spike rate for those units within the ensemble was higher when the ensemble was active than when it was inactive (**Figure 3C**, green, two-sided Wilcoxon rank sum test, $Z = 20.9$, $D = 0.8$, power = 0.99, $p < 0.001$). We also assessed the average spike rate when all single units within an ensemble were merged into a single multi-unit spike train (**Figure 3D**). Again, spike rate within the ensemble was higher during epochs of ensemble activation (**Figure 3D**, green, two-sided Wilcoxon rank sum test, $Z = 14.7$, $D = 2.6$, power = 0.99, $p < 0.001$). On the other hand, when an ensemble was inactive, multi-unit activity outside of the ensemble was relatively higher (**Figure 3D**, orange, two-sided Wilcoxon rank sum test, $Z = 6.8$, $D = 0.8$, power = 0.99, $p < 0.001$). Presumably, this is due to those units spiking as members of other ensembles during those epochs. These findings confirm what is shown in **Figure 1B**, namely that LC ensemble activations are associated with an increase in spike rate of only the single units in that ensemble. Most importantly, these results show that the firing strength can vary considerably across LC ensembles, which could potentially correlate with their relation to cortical state.

Next, we assessed how ensemble activation related to the tendency of LC single units to fire in bursts. We defined a burst as an occurrence of 2 or more consecutive spikes with an inter-spike interval of less than 80 msec. We chose 80 msec based on the physiological definition used in prior work on the LC, as well as on dopamine neurons (Grace and Bunney, 1984; Tung et al., 1989).

Low inter-spike intervals are important from the point of view of the physiological downstream effect LC neuronal spiking, as it is known that when LC neurons send a few spikes in a short ISI, K+ leak causes axonal depolarization so that later spikes in the burst are conducted faster (Aston-Jones et al., 1985); therefore, the ISI at the sites of norepinephrine release may be even shorter than measured at the soma. Direct electrical stimulation of the LC at burst frequencies increases norepinephrine release (Florin-Lechner et al., 1996). Therefore, brief ISIs during ensemble activation could play a key role in the contribution of LC ensemble activations to cortical state. Within ensembles, single units tended to burst more during the ensemble active times than during inactive times (**Figure 3C**, two-sided Wilcoxon rank sum test, $Z = 15.2$, $D = 0.5$, power = 0.99, $p < 0.0001$). We found that single units within an ensemble, when merged into a single multi-unit spike train, also burst more often when the ensemble was active than when the ensemble was in an inactive state (**Figure 3C**, two-sided Wilcoxon rank sum test, $Z = 13.1$, $D = 1.0$, power = 0.99, $p < 0.0001$). These results demonstrate than LC ensemble activation is associated with increased burst firing of the units in the ensemble, which in turn suggest a strong downstream effect of LC ensembles on their (forebrain) projection targets.

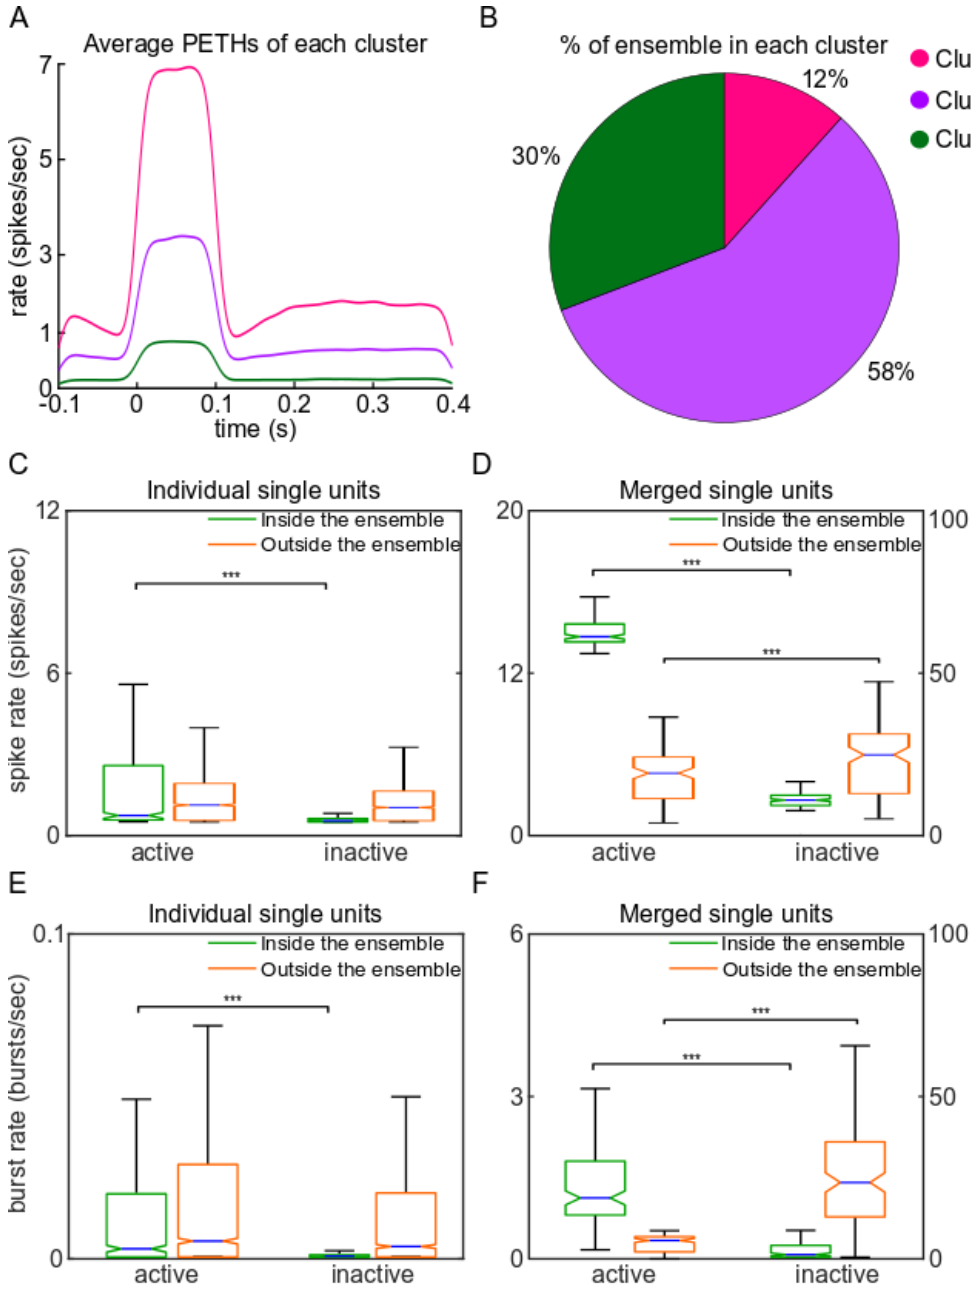

Figure 3. LC ensembles are characterized by different magnitudes of change in spike rate and an increase in burst firing. (A) Average PETHs of the ensembles in the same cluster. The zero time on the x-axis is the ensemble active time. The PETHs of all ensembles were grouped into 3 clusters that increased their firing rate to different degrees. **(B)** The pie chart illustrates the percentage of ensembles in each PETH cluster. Most ensembles had a medium (purple) or low (green) magnitude increase in single unit spike rate. **(C, D)** The box plots show the distribution of the spike rates for the single units inside the ensemble (green) and outside the ensemble (orange). The result is shown separately for individual single units (C) and the spike trains merged across single units (D). The spike rate was calculated as the average of all ensemble activation events combined across all single units in the ensemble (i.e., in Figure 1B and 5A spike rasters, all events

of different colors were averaged). Spike rate increased when the ensemble was in an active state for both single units and multi-unit activity. Additionally, when the ensemble was inactive, multi-unit activity outside of the ensemble was higher presumably due to units spiking in other ensembles. **(E, F)** 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 individual single units (E) and the spike trains merged across single units (F). Single units inside the ensemble burst more frequently during ensemble active times. The same result was found for merged spike trains of all single units in the ensemble (F, green). However, merging the spike trains of single units outside of the ensemble revealed increased bursting when the ensemble was not active. Again, this difference is presumably due to the units outside of a selected ensemble being active in other ensembles when the selected ensemble in offline.

LC ensembles show signs of self-inhibition and limited lateral-inhibition

Our analyses have shown that LC ensembles are activated briefly with long pauses between an activation and the next one (**Figure 1B**). Such pauses may maintain independent activation times between ensembles, so that different ensembles can activate at different times and produce ensemble-specific cortical states. Pauses between LC ensemble activations may be generated by local inhibitory mechanisms. Local noradrenergic inhibition is a prevalent determinant of the spiking patterns of individual LC neurons via self-inhibitory and lateral-inhibition neuronal circuit motifs (Aghajanian et al., 1977; Ennis and Aston-Jones, 1986). LC neurons are large (25 um soma) and densely packed with numerous close proximity dendrites (Swanson, 1976; SHIMIZU et al., 1978; Grzanna and Molliver, 1980), which are the site of alpha-2 receptors that can mediate noradrenergic self-inhibition and lateral-inhibition (Lee et al., 1998; Huang et al., 2007). Thus, the activation of multiple LC neurons and the volume transmission of local release 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 initially excites LC neurons and the resultant local norepinephrine release and its volume transmission causes inhibition of all recorded neurons around 200 um from the stimulation site (Marzo et al., 2014). We predicted that LC ensemble activations, which involve synchronous activation of multiple LC neurons, would be associated with a similar spread of lateral inhibition across many single units. Such synchronous release of norepinephrine by the neurons in an ensemble should also serve to self-inhibit the ensemble.

The NMF method for 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. We assessed self-inhibition of LC ensembles by examining LC ensemble activation timing auto-correlograms (**Figure 4A**). We assessed lateral-inhibition between LC ensemble-pairs by measuring the cross-correlograms between their activation times (**Figure 4B**). We found that a trough in the auto-correlogram, which indicates self-inhibition, occurred in the majority of the ensembles (90 out of 146 ensembles, 62%). Of these 90 ensembles with signs of self-inhibition, the inhibition lasted less than 300 msec and the spiking was most inhibited at 100 msec after ensemble activation (**Figure 4E**). When we considered lateral-inhibition between pairs of LC ensembles, we observed 44% of ensemble-pairs (out of 790 total) had an inhibitory interaction. The histogram showing the timing of significant lateral-inhibitory interactions between ensemble-pairs has a peak at ±300 msec (**Figure 4F**). Overall, these analyses demonstrate some similarities between LC functional motifs for single units and ensembles. Specifically, we show that LC ensembles tend to inhibit themselves. Moreover, we show that some ensembles laterally-inhibit other ensembles. However, the activation of an LC ensemble does not cause a global "halo" of surrounding inhibition across the LC given that only 44% of ensemble-pairs showed signs of lateral-inhibition. These inhibitory mechanisms could help produce the sparse activations of LC ensembles, replete with pauses, such that ensembles activate largely independently from each other.

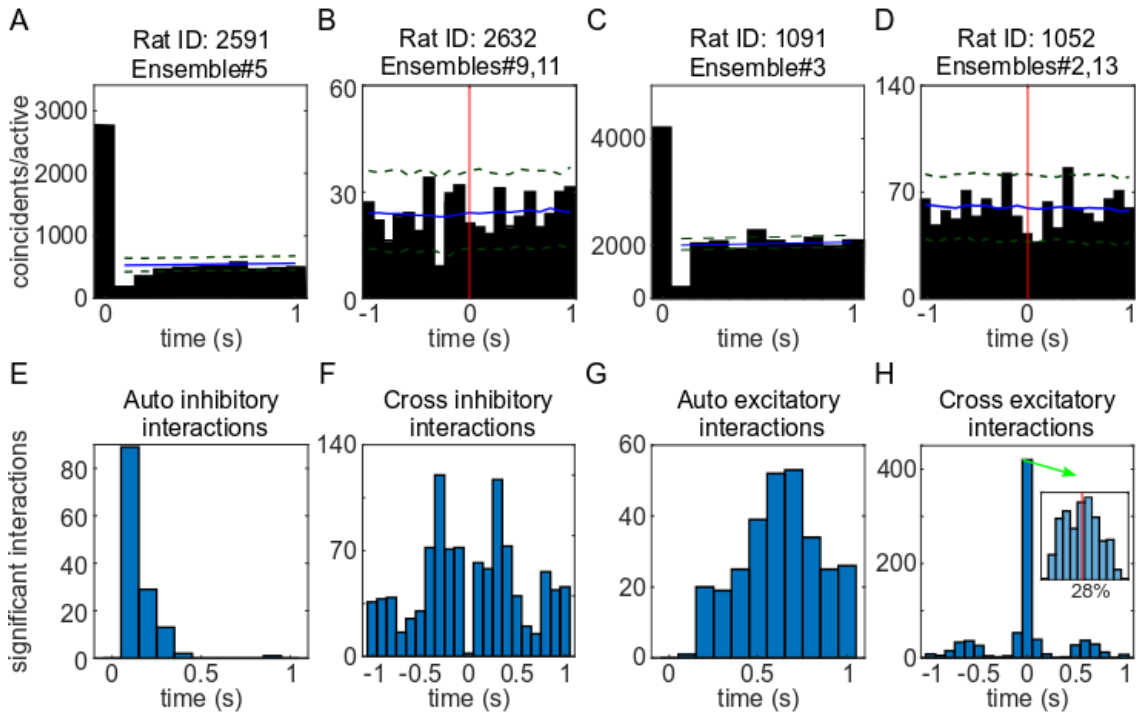

Figure 4. LC ensembles exhibited self-inhibition, lateral-inhibition, and patterned self-excitation and co-excitation. (A-D) Four examples of self-interactions (A - inhibitory, C - excitatory) and cross-interactions (B - inhibitory, D - excitatory). Significant excitatory or inhibitory interactions were 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 1000 surrogate data sets constructed by jittering ensemble active times. The solid blue line shows the average of the surrogate correlograms. **(E)** Histogram showing the number of significant self-inhibitions during different time bins. The plot shows that self-inhibition in almost all cases (98% of all auto-correlogram time points across all ensembles) lasts less than 300 msec. **(F)** Histogram showing the number of significant lateral-inhibition times for all ensemble-pairs that exhibited significant lateral-inhibition. The histogram shows a peak at ±300 msec. **(G)** Histogram showing the number of significant self-excitations during different time bins. Self-excitation happens after 300 msec in 73% of the ensembles. **(H)** Histogram showing the number of significant ensemble-pair coactivation during each time bin reveals an initial peak at time 0 and another peak around ±600 msec. The inner plot shows the histogram of synchrony index values between the ensemble-pairs that had significant coactivation at time 0. The average synchrony index is 28%, which indicates low zero-lag synchrony, even among ensemble-pairs with a significant zero-lag peak in the cross-correlogram.

LC ensemble pairs exhibit time-patterned excitations

Another important piece of information for understanding the temporal organization of LC ensemble activity is how often an LC ensemble self-activates and how often it activates in time with other LC ensembles. Although the LC has no intrinsic excitatory neurotransmitters, it does receive numerous extrinsic sources of excitatory input (for review, see Totah et al., 2018b). Thus, 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) that systematically pattern the activation times of LC ensembles. We examined the temporal pattern of LC ensemble activation events by calculating LC ensemble auto-correlograms (to study self-excitation) and LC ensemble-pair cross-correlograms to study coactivation properties of LC ensembles. In LC ensemble activation auto-correlograms, we observed self-excitation occurring after 300 msec in 73% of the 146 ensembles (**Figure 4G**). Among 790 ensemble-pairs, we observed 64% had excitatory interactions which peaked around time zero and again around ±600 msec (**Figure 4H**), indicating some degree of temporal organization in the sequency of activation times of different ensembles. We quantified the amount of zero-lag synchrony between ensembles using a synchronization index (see Methods). The average synchrony between the ensembles with

significant coactivation at time 0 was around 28% (**Figure 4H**, inset). The low value of the synchronization index shows that ensembles that have synchronous coactivation (i.e., significant zero-lag values in the cross-correlogram) are, on average, activated together only in a limited proportion (on average, 28%) of instances of activations. Thus, contrary to the traditional hypothesis that LC fires *en masse* with a high level of population synchrony, ensembles are rarely firing in zero-lag synchrony. Overall, excitation of LC ensembles occurs in diverse patterns with a preference for partly rhythmic excitatory interactions with a time lag of ~600 msec and rarer bouts of ensemble-pair synchrony.

Diverse cortical states are associated with activation of different LC ensembles

Our analyses thus far have demonstrated that the LC has unit type-specific ensembles which are activated largely at different times and, on rare instances, coactivated. We turned to our other central objective, which was to examine the relationships between brain state and LC ensemble dynamics. Prior work has demonstrated that LC neurons, when synchronously activated, cause a specific change in cortical oscillatory state during either anesthesia or sleep that is characterized by decreased low frequency spectral power and increased higher frequency spectral power in the mean extracellular field potential (Eschenko et al., 2011; Marzo et al., 2014; Safaai et al., 2015). This cortical state has been termed the "activated state." It is possible, however, that activation of different LC ensembles could be associated with other brain states. LC neurons have localized projections to the forebrain and release a range of neurotransmitters (Totah et al., 2018b; Chandler et al., 2019); therefore, LC ensembles that project to different forebrain neuronal networks could affect how those neuronal networks self-organize brain states.

We tested the hypothesis that activation of LC ensembles was accompanied by the activated cortical state. Such a finding would be wholly consistent with prior studies during urethane-anesthesia or during sleep, which examined cortical spectral power after spontaneous increases in LC multi-unit spiking or after LC stimulation (Carter et al., 2010; Marzo et al., 2014; Neves et al., 2018; Hayat et al., 2020). Here, we examined changes in cortical area 24a local field potential (LFP) power triggered on LC ensemble activation times. For each instance of ensemble activation, we calculated the LFP spectrogram modulation in a 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 4**) and durations of activation and inactivation (**Figure 2A**) show that it is unlikely that ensembles were coactive during this window. Therefore, this window ensured that changes in the cortical LFP spectrum were related to activation of an individual ensemble. We averaged the spectral modulations for each ensemble over all instances its activation (N = 89 ensembles considered for this analysis because cortical LFP was recorded in 9 out of 15 rats). Visual inspection of the LC ensemble activation-triggered spectra revealed diverse cortical states depending on which ensemble was activated (**Supplementary Figure 4)**.

We thus assessed whether LC ensemble activation gave rise to multiple types of cortical states and, if so, which were the most predominant cortical states associated with LC ensemble activation. To do so, we clustered the spectral modulations associated with each of the 89 LC ensembles. If all LC ensembles were associated with the cortical activated state, then a single cluster would be observed. However, we found 4 predominant types of spectra in the clustering analysis out of which only one can be described as the activated state (i.e., decreased low frequency spectral power and increased high frequency spectral power). We chose 4 clusters by first varying the putative number of clusters from 1 to 22 and then choosing the number of clusters as the elbow of the curve where error dropped below 5% such that using more clusters would have given a much diminished return in terms of additional quality of data description (for details, see Methods and **Supplementary Figure 5A**). A silhouette analysis was used to confirm consistency of clustering across different distance functions (**Supplementary Figure 5B**). The average spectrum of each spectral cluster is shown in the top row of **Figure 5A** (non-significant modulations are white) and the lower rows show the activity of a single example ensemble selected from each cluster and that ensemble's activation-triggered cortical spectrum. Across all plots, the time of LC ensemble activation (t=0 sec) is indicated by a red line. Our results clearly show that activation of different sub-sets of ensembles is associated with different cortical states beyond the canonical activated state.

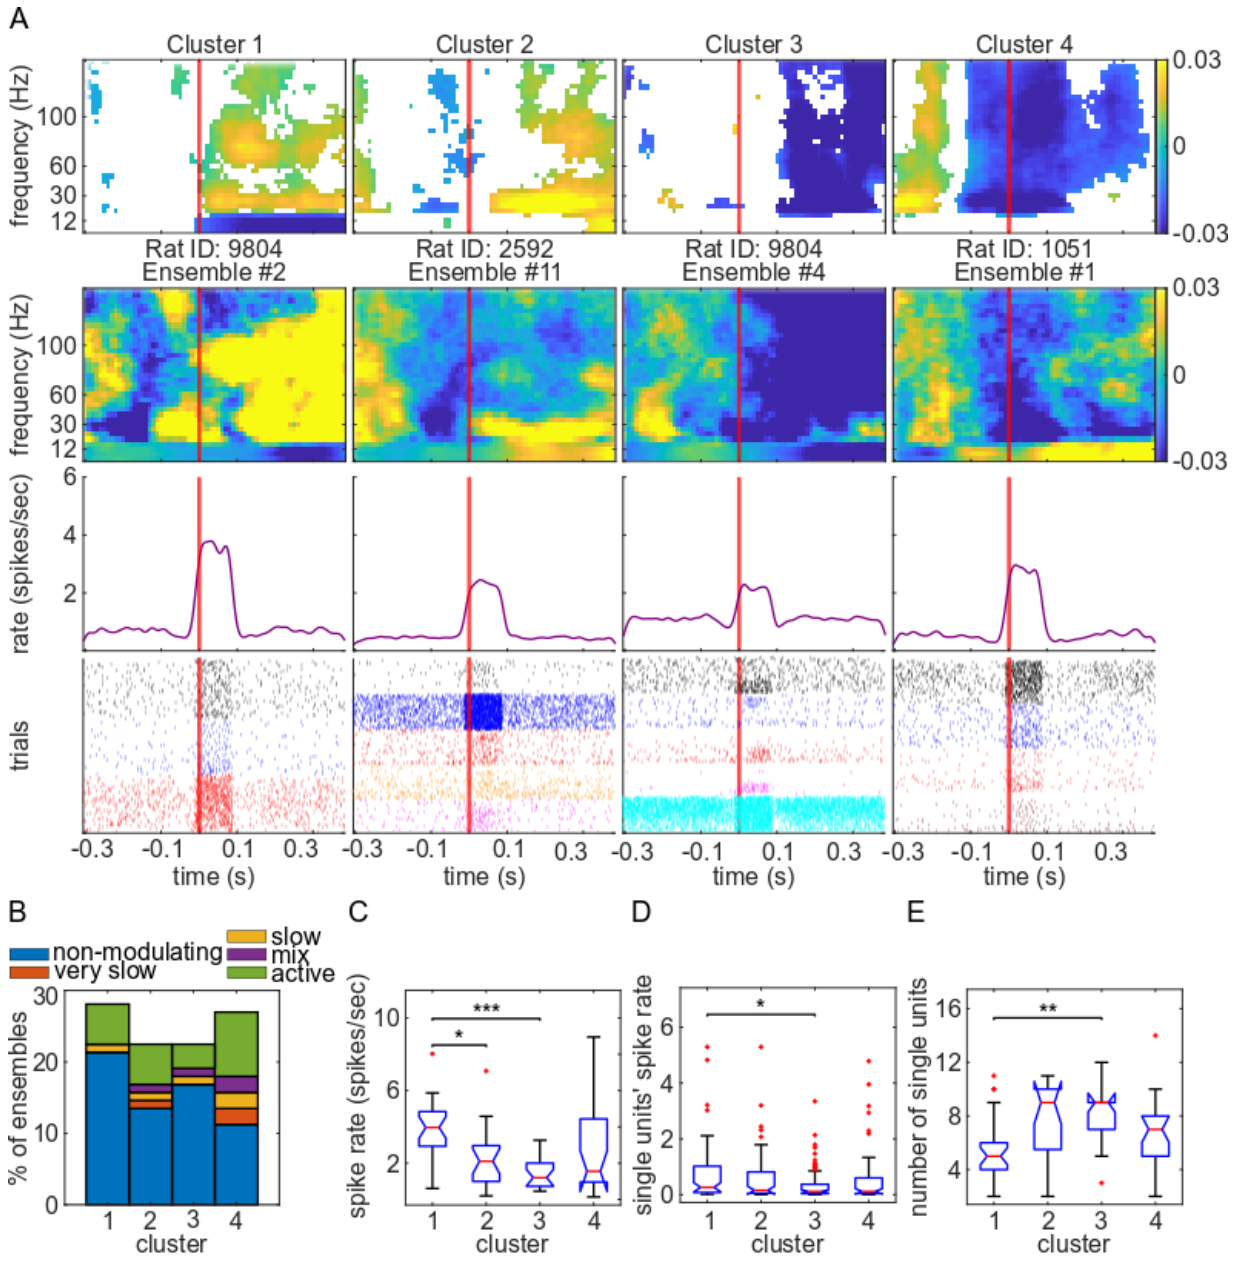

Figure 5. Activation of different LC ensembles are associated with diverse changes in cortical LFP power spectra. (A) LFP power spectra were triggered on LC ensemble activation times. The resulting spectra were clustered into 4 types, which are shown as 4 columns. The top row shows the average spectrogram across all spectra in each cluster. Only significant modulations (yellow – increase, blue – decrease) are shown; non-significant 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 of the 4 spectral clusters and the ensemble activation-triggered spectrogram for that ensemble. The example spectra show both significant and non-significant values. **(B)** A histogram of the percentage of the ensembles in each cluster that were preferentially active during a specific ongoing cortical state. The blue bar shows ensembles with no preferred state. Most

ensembles are not preferentially active during a specific ongoing cortical state, but the remaining ensembles were preferentially engaged during the activated state. **(C)** The box plots show the distributions of maximal spike rates of the ensembles' PETHs in each spectral cluster. There was a significant difference in spike 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 difference was found only between spectral clusters 1 and 3.

The power spectra were spatiotemporally diverse across the 4 clusters. Ensembles that fell into the first type of spectrum (Cluster 1 in **Figure 5A**, top row) were associated with the activated state. This type of spectra was associated with activation of 28% of the 89 ensembles. The second spectral type (Cluster 2) was associated with activation of 22.5% of LC ensembles and was characterized by an increase in middle to high frequency components of the LFP whereas low frequencies did not change. The third type of spectral modulation (Cluster 3) opposed the direction of the first two spectral types, in that the middle to high frequency components of the LFP were decreased. This spectral pattern was associated with 22.5% of the ensembles. In all 3 of the aforementioned spectral types, the change in cortical state took place after LC ensembles activated. However, the last type of spectrum (Cluster 4) was associated with a change in cortical state that began before LC ensemble activation, namely a decrease in high frequency spectral power. Overall, different LC ensembles are associated with spectrotemporally diverse cortical states.

One explanation for the diversity of LC ensemble-specific cortical states is that specific ensembles are active only during particular types of cortical states. Under urethane anesthesia, cortical LFP can be characterized by relatively long durations (multiple seconds to minutes or hours) of predominantly very slow (<1 Hz) / slow (<4 Hz) oscillations or by the "activated" state (Clement et al., 2008). For instance, it is possible that ensemble activations associated with a type 1 spectral cluster (**Figure 5A**) could be observed only during the activated state and thus represent a strengthening of that state (i.e., a further reduction in low frequency spectral power and a potentiation of high frequency spectral power). Here, we examined this possibility by assigning ongoing states to long epochs (7.5 sec duration) of cortical activity and assessing the preferred ongoing state (or lack thereof) of the LC ensembles grouped into different spectral cluster types. Each 7.5 sec window of cortical LFP was assigned one of four possible cortical states based on

prior definitions (see Methods and (Totah et al., 2018a): very slow, slow, mixture of slow and activated, or activated. A Bayesian procedure was then used to assess whether ensembles tended to be active in a preferred ongoing cortical state or not. Overall, we found that most ensembles (63%) were not preferentially active during any specific ongoing cortical state (**Figure 5B**). This finding shows, for instance, that a type 1 cortical power spectrum associated with the activation of a particular LC ensemble was just as likely to occur during an epoch of ongoing cortical slow oscillations as during an epoch of ongoing cortical activation. However, approximately 24% of the ensembles were preferentially active during the activated state. Therefore, all 4 types of cortical power spectra shown in **Figure 5A** could be observed during longer epochs (7.5 sec) of cortical activation. For example, activation of an LC ensemble associated with a type 1 spectral cluster occurring during the activated state could be interpreted as a strengthening of the activated state. On the other hand, a type 2 spectral cluster could be interpreted as a modification of the ongoing activated state, whereby power increases in only the middle to high frequencies. 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). This analysis shows that the diverse set of cortical states associated with activation of different LC ensembles is not due to specific ensembles being active during only particular ongoing cortical states.

Another potential explanation for the distinct cortical states associated with different LC ensembles is that the firing strength of the ensemble is related to the cortical state. For instance, the canonical activated state (type 1 cortical power spectrum) might only occur when an LC ensemble reaches a certain threshold spike rate. We showed that different ensembles have different firing strengths (**Figure 3A**) and these differences might predict the relationship between the LC ensemble and cortical state. We assessed whether there was a systematic difference in the strength of ensemble population spike rate across the 4 cortical state clusters shown in **Figure 5A**. The population spike rate was calculated as the average of all ensemble activation events combined across all single units in the ensemble (i.e., in the spike rasters shown in **Figure 1B** and **Figure 5A**, all events of different colors were averaged). The peak spike rate of the resulting PETH was used to characterize the firing strength of the population in each ensemble. We found that the median population spike rate across ensembles in each cortical spectral cluster differed across clusters (Kruskal-Wallis test, $p = 0.0003$, $\omega^2 = 0.9633$, $\chi^2 = 18.82$), but post-hoc tests showed that

only cluster 1 was different from clusters 2 and 3; therefore, there was no systematic relationship between cortical spectral cluster type and population spike rate (**Figure 5C**). We also examined the peak spike rate of the single units in each ensemble. For this analysis, the spike rates around ensemble activation events were first averaged for each single unit separately (i.e., in the spike rasters shown in **Figure 1B** and **Figure 5A**, all events of the same color were first averaged). The peak spike rate of the PETH of each single unit was averaged across units to obtain a measure of single unit firing strength. The median spike rate across all ensembles in each cortical spectral cluster type again differed across clusters (Kruskal-Wallis test, $p = 0.0334$, $\omega^2 = 0.9871$, $\chi^2 = 8.71$). The result was similar to that obtained by examination of the population spike rate, in that the single unit firing rate differed only between clusters 1 and 3 (**Figure 5D**).

We examined a final factor that could predict how different LC ensembles are associated with particular cortical states. Specifically, the size of the ensemble (i.e., the number of single units within the ensemble) might systematically vary with the cortical spectral cluster type. For instance, a type 1 cluster might only be observed when ensembles of a particular size are activated. In order to assess this relationship, we calculated the median number of units across ensembles in each cortical state cluster. Ensemble size differed across clusters (Kruskal-Wallis test, $p = 0.0029$, $\omega^2 = 0.9608$, $\chi^2 = 13.97$), but only between clusters 1 and 3 (**Figure 5E**). These results demonstrate that, while cluster 1 and 3 differ, there is no systematic relationship between the size of an ensemble and cortical state. Overall, our results demonstrate that cortical state depends on which specific ensembles are active, rather than simply an overall increase in the number of active single units or their firing strength.

Activating a larger pool of LC ensembles results in a more homogeneous cortical state

Our data clearly demonstrate a relationship between individual LC ensembles and distinct cortical states with heterogeneous spectrotemporal properties. This finding stands in marked contrast to the single activated state evoked by direct stimulation of the LC, which activates most of the neurons synchronously (Marzo et al., 2014). Therefore, we predicted that when LC ensembles are coactive (i.e., more of the LC neurons are activated simultaneously and approaching stimulation-evoked whole-LC activation), the associated cortical state should become more homogenous. In order to test this prediction, we took advantage of our observation that pairs of LC ensembles can

sometimes become coactive (28% of the time among 64% of the ensemble-pairs, **Figure 4H**). We assessed the cortical LFP spectra, as in **Figure 5**, but triggering cortical spectrograms only on coactivation times of ensemble pairs (**Figure 6**). A total of 199 ensemble-pairs had a significant zero lag cross-correlogram indicating coactivation. In contrast with the four heterogenous cortical states observed during activation of individual LC ensembles (**Figure 5A**), k-means clustering now revealed only two types of cortical power spectra (**Figure 6**). One cluster is the canonical activated brain state that is expected based on prior studies of LC activity (cluster 2, 103 of 199 ensemble-pairs) and the other cluster is a homogenous decrease in spectral power (96 ensemble-pairs in cluster 1). Therefore, when multiple LC ensembles are coactive such that the LC population activity is closer to whole-LC activation, the modulation of cortical state is more homogenous.

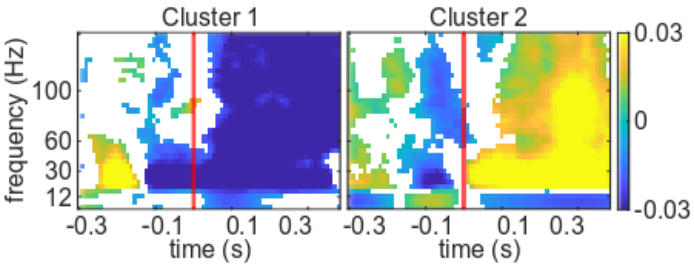

Figure 6. Synchronous coactivations of LC ensemble-pairs are associated with more homogeneous changes in cortical LFP power spectra. The LFP power spectra were triggered on the coactivation times of LC ensemble-pairs that had significant zero-lag cross-correlations (**Figure 4H**). The resulting spectra clustered into 2 spectral types. Each plot shows the average spectrogram across all ensemble-pairs associated with each type of spectrum. Only significant changes in the power spectrum are plotted in color; non-significant modulations are white.

Discussion

Cortical states can vary over a wide range and have been shown to be in a tight relationship with many functions that are relevant to psychiatric disorders, such as arousal level, perceptual ability, cognitive task engagement, and reaction times. It is thus no surprise that there have been long-standing efforts to understand the neural factors contributing to cortical state fluctuations (Harris and Thiele, 2011; McGinley et al., 2015; McCormick et al., 2020). While the LC has been long implicated in having a role in determining cortical state maintenance and transitions, LC neuronal

population activity has always been thought and shown to produce a single activated cortical state (similar to that observed in sleep-wake transitions) by presumably collective firing of the LC neurons (Carter et al., 2010; Marzo et al., 2014; Hayat et al., 2020). However, the LC is of course involved in many other functions beyond the sleep-wake transition, which begs the question of whether LC population activity, perhaps through more nuanced dynamics than *en masse* firing, could be involved in the control of a diversity of brain states beyond the activated state.

Here, we use NMF to decompose LC population activity into individual patterns of coactive neurons over time. Our analyses show that the population activity consists of a nuanced sequence of different ensembles activating at largely different times and only rarely in synchrony. We characterized the physiological features of ensembles and found variations in firing strength and size of the neuronal population, as well as increased burst firing during ensemble activation. Ensembles tended to be made of units of the same type, with individual units mostly participating in one ensemble, but sometimes in more than one ensemble. Analysis of the temporal dynamics of ensemble activity revealed that LC ensembles are self-inhibitory and also laterally-inhibit one another, which may contribute to the rarity of simultaneous activation of multiple ensembles.

We addressed our central question about the LC relationship with cortical state by triggering cortical power spectra on LC ensemble activation times. Our analysis revealed four types of LC ensemble-associated states, which were spectrotemporally diverse. 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. Importantly, coactivation of LC ensemble-pairs, which approaches closer to the whole-LC activation caused by LC stimulation, was associated with only two cortical states, including an activated one. Our results are in line with stimulation studies showing that nearly whole-LC activation produces the activated state exclusively (Carter et al., 2010; Marzo et al., 2014; Hayat et al., 2020). However, contrary to the current framework that the LC produces the activated state, our findings suggest that discrete LC ensembles are spontaneously activated

612 with a rich, largely non-synchronous dynamic and that different ensembles associate with
613 unexpectedly diverse cortical states.

614 *Spontaneous activation of LC ensembles at different times may be enforced by local noradrenergic*
615 *inhibition*

616 We observed that LC ensembles were activated at largely different times. Such activation patterns
617 must rely on afferents to the LC and intra-LC neurotransmitters. Our analyses of auto- and cross-
618 correlations of ensemble activation times suggest that intra-LC norepinephrine volume
619 neurotransmission may be involved in maintaining the independence between LC ensembles and
620 structuring their activation timing by regulating epochs of ensemble silence. We found that self-
621 inhibition was maximal 100 msec after ensemble activation and lasted a few hundred msec. This
622 duration is similar to spontaneous self-inhibition of LC single units (Ennis and Aston-Jones, 1986).
623 It is also consistent with the duration of self-inhibition after a single current pulse evokes an
624 increase in spontaneous firing (Marzo et al., 2014). The synchronized activation of all single units
625 in an ensemble likely causes a post-activation inhibition that is also synchronized across the single
626 units in that ensemble. Thus, noradrenergic self-inhibition by LC ensembles could potentially
627 rapidly curtail an ensemble's activity after it has fired, which would reduce synchrony across
628 ensembles and promote independence between ensembles.

629 Lateral inhibition may offer a similar constraint on LC ensemble activation due to the intra-LC
630 volume transmission of norepinephrine released during neuronal activity, which would stimulate
631 alpha-2 noradrenergic receptors on neurons in other ensembles. We observed lateral inhibition
632 between LC ensembles, which was apparent as decreased coincidental ensemble-pair activations
633 in cross-correlograms. Although activation of multiple LC neurons in an ensemble might be
634 expected to produce a "halo" of surrounding inhibition across the LC, we observed lateral
635 inhibition among only 44% of ensemble-pairs. Overall, these analyses demonstrate some similar
636 noradrenergic inhibitory motifs that apply to both single LC neurons and LC neuronal ensembles
637 and may promote activation of ensembles at distinct times.

638 *Potential causes of the diversity in cortical state during activation of LC ensembles*

The physiological causes of LC ensemble-specific cortical states are unknown, but two potential factors for future study are the diversity of each ensemble's neurochemical make-up and/or its projection profile. Given that the region in which we assessed cortical state (area 24a) receives projections from approximately 61 to 65% of LC neurons in the rat (Chandler et al., 2013, 2014), it seems likely that most of the ensembles project to area 24a and they should, therefore, produce a similar state change. Our finding to the contrary could be explained by the possibility that the neurochemical make-up of the LC neurons differs across ensembles and results in cortical state diversity. Another possibility is that, in spite of most ensembles presumably sharing area 24a as a projection target, it is the other targets that are potentially not shared across ensembles, which leads to LC ensemble-specific cortical states in 24a. According to this network perspective, LC ensembles associated with different brain states have divergent axon collaterals which enable the ensembles to modulate distinct neuronal networks (i.e., sub-sets of brain regions). The neuromodulation of different networks changes the self-organization of cortical states.

It is worth noting here that we observed patterns of LC burst firing during ensemble activation that suggest norepinephrine release could be temporally-coordinated across multiple brain targets, which could allow the LC neurons in a specific ensemble to modulate a specific multi-region neuronal network. We found that LC ensemble activity was associated with burst firing (i.e., <80 msec inter-spike interval) both at a population level across units in the ensemble and for spikes of individual single units in the ensemble. Population level bursts may be relevant to tightly-timed norepinephrine release in multiple brain regions innervated by the units in the ensemble. On the other hand, increased bursting by single units during ensemble activation may promote increased norepinephrine release from single neurons by altering LC axon conduction velocities and neurotransmitter release probability (Aston-Jones et al., 1980; Florin-Lechner et al., 1996). Overall, the burst firing patterns we observed could serve to potentiate the amount of norepinephrine release, while at the same time temporally-coordinating that "boost" in norepinephrine release across multiple targets of the neurons in the ensemble. Population level bursts and single unit bursts during ensemble activation may be another physiological characteristic that influences the relationship between specific LC ensembles and various cortical states.

Implications for understanding the role of LC in brain state and behavioral transitions

Our findings open up the intriguing possibility that the brain state during specific behavioral events may depend on which LC ensemble(s) activate. New tools that finally enable large-scale recordings of many LC single units simultaneously in the awake organism could reveal such a relationship. Although we observed rare coactivation of LC ensembles (28% of the time on average among 64% of ensemble-pairs) under anesthetized conditions, it is possible that the LC ensemble activation patterns differ in the awake organism. For instance, among cholinergic neurons, highly-transient synchronous ensemble activity has been observed at the time of behavioral transitions (i.e., locomotion onset) in mice (Howe et al., 2019). However, given that anesthesia promotes synchronous firing among neurons, it is likely that LC ensembles are independently-active in the non-anesthetized organism too. Overall, our analyses demonstrate for the first time that LC population activity is constructed from independently active ensembles and that the canonical activated state associated with *en masse* collective activation of LC neurons is only one of many brain states that can occur depending on which particular LC ensemble is active.

Materials and Methods

Data collection: Recording procedure 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 during the active period of the rats, which were housed on a light cycle of 08:00 to 20:00 darkness. A sub-set of the data were collected from rats used in a prior study and typical histological sections are shown in that work (Totah et al., 2018a).

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 (Totah et al., 2018a). 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 electrophysiological criteria (see prior work for a detailed description (Totah et al., 2018a). At the end of the recording, we administered clonidine (0.05 mg/kg) i.p. (Sigma-Aldrich, product identification: C7897) to confirm cessation of noradrenergic neuron spiking. We also verified LC targeting in most experiments using histological examination of coronal sections (50 um thick) that were stained for Cresyl violet or a DAB and horse radish peroxidase reaction with hydrogen peroxide to visualize an antibody against tyrosine hydroxylase, as shown in prior work (Totah et al., 2018a).

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

NMF decomposition of population spike trains into coactive ensembles

We used non-negative matrix factorization (NMF) (Lee and Seung, 1999) to decompose a matrix of the spike counts of all simultaneously recorded single units across time intervals. NMF linearly decomposes the matrix of the spike counts of the population of single units at each time interval as a sum across a set of non-negative basis functions (modules) using non-negative coefficients (Lee and Seung, 1999; Onken et al., 2016; Williams et al., 2018). The non-negativity constraint is useful for obtaining sparse representations and it is particular suitable for decomposing population spike count at different time intervals, which are always non-negative. Previous work has shown that the NMF of population spike trains provides a robust decomposition whose basis functions can be biologically interpreted as a set of the firing patterns of the single units that are coactive (i.e., an ensemble) and the coefficients quantify the relative strength of recruitment of each ensemble firing pattern at any given time (Onken et al., 2016).

We employed an NMF decomposition that we have previously termed "space only NMF" because it decomposes the population firing patterns across single units at each time interval (Onken et al., 2016):

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

$R \in \mathbb{Z}_+^{T \times N}$ is the data matrix containing the spike counts of each of N single units binned into T time bins (with t being the index of each time bin). $H \in \mathbb{R}_+^{K \times N}$ is the matrix containing the basis function, which has K spatial modules. Each module captures a different pattern of coactivity of the single units and can, therefore, be used to identify which neurons are active together and thus form ensembles. $W \in \mathbb{R}_+^{T \times K}$ is the matrix containing the activation coefficients that describe the strength of recruitment of each module (and thus of each ensemble of coactive neurons) at each time interval. The residuals express the error in the reconstruction of the original population spike train matrix. We computed the decomposition using the multiplicative update rule to minimize the Frobenius norm between the original and the reconstructed data (Lee and Seung, 1999). Note that the use of the Frobenius norm assumes that the residuals have a Gaussian white noise structure.

One free parameter of the analysis is the temporal resolution of the time binning, ΔT. We binned spike counts at ΔT = 100 msec. The time resolution was selected based on our previous work reporting that pairs of LC single units are predominantly synchronized on a timescale of approximately 100 msec or less (Totah et al., 2018a). We also used ranges of ΔT from several tens of msec to a few hundreds of msec and found that shorter bins (<=20 msec) and longer bins (> 1s), which our prior work suggests would be outside the range of LC single unit synchrony, tend to artificially identify either many modules each containing only one single active unit or one large ensemble containing all single units, respectively.

The second free parameter of the NMF analysis is the number of different modules,K, which were chosen for computing the decomposition. Following established procedures (Onken et al., 2016; Williams et al., 2018), we chose K for each rat by computing the amount of the variance explained by the decomposition when varying K from its minimum possible value (one) to its maximum possible value (the number of simultaneously-recorded single units). An elbow in this plot indicates a point of diminishing returns for including more modules. We thus chose the number of

modules as the smallest K in the elbow region of this curve for which the decomposition reconstructed at least 60% of the variance of the original spike train data. Given that the NMF decomposition may have local maxima in the variance explained (or equivalently local minima in error reconstruction), after selecting K , we repeated the decomposition five times using this K and used randomly chosen initialization conditions on each repetition. The selected K was used if all solutions had a high degree of stability across these five random initializations. Stability was assessed by checking the repeatability of clustering in comparison to randomly assigning single units to ensembles. The degree of stability was computed as follows. We hard clustered the data to assign each single unit to one and only one ensemble by dividing each column of H by its maximum and removing the values below 1. From these data we then measured the stability across the five decompositions using the Rand Index (Rand, 1971). We compared the average of the Rand index for each animal with 100 repetitions of the five random clustering. The average Rand Index 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. Among those random initializations, the final decomposition reported in the analyses was chosen as the one leading to the maximum variance explained.

The modules detected by NMF provide a pattern of coactivation of different single units and the activation coefficients measure the strength of recruitment of each module at any given time. From these data, we used a threshold-crossing of the coefficients to define when ensembles were active and which single units were active in the ensemble. In order to perform the thresholding, we first normalized the columns of H to the minimum and maximum and then set a threshold based on the distribution of coefficients. Single units within a module were defined as an ensemble of coactive single units if their corresponding element of H crossed the 95^{th} percentile threshold of the distribution of coefficient values for that rat. 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 activation coefficients representing the strength of recruitment of each module at any given time interval. We thresholded these continuous values into binary values using the same method explained above for the spatial modules. The binary

783 version of the matrix, W, hereafter referred to as "activation coefficient matrix," was used to
784 determine whether an ensemble is active or not in each time bin.

785 *The evaluation of physical clustering of ensembles according to location on the recording array*

786 To assess whether single units within an ensemble tended to cluster on the recording array, we
787 measured the pairwise distance between the units within each ensemble. The location of each unit
788 was assigned to the electrodes on which the maximal waveform was recorded. Euclidian pairwise
789 distances of the units inside each ensemble were calculated.

790 *The assignment of single unit types in the ensembles*

791 Single unit type was defined by waveform duration, as in prior work (Totah et al., 2018a). We
792 determined if single units of the same type were more likely than chance to belong to an ensemble
793 by computing the exact probability of having ensembles of the same single unit type under the null
794 hypothesis of random assignment. These probabilities were computed by the means of repetition
795 of random sampling (assembling) without replacement. The number of units in the sample was
796 fixed to the number of single units in the ensemble. The number of repetitions for each rat was the
797 number of ensembles that were empirically found by NMF to consist of only one type of single
798 unit.

799 *Calculation of cross-correlograms between pairs of ensembles*

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

806 The degree of synchrony between ensemble-pairs that had a significant excitatory interaction at
807 time 0 in the cross-correlogram was measured using a synchrony index:

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

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.

Calculation of ensemble auto-correlograms

Auto-correlograms were calculated in a 1000 msec time window using a 100 msec time bin. The significance of inhibition was assessed with the same procedure used for cross-correlograms (see above) that compared the observed auto-correlogram against 1000 surrogate auto-correlograms.

Peri-event time histograms of single unit activity during ensemble activation

The Peri-Event Time Histogram (PETH) of 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 spike rate 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 different "trials". PETHs were smoothed with a Gaussian kernel (10 msec width). The PETH for each ensemble was obtained by averaging PETHs across all single units that were active in 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 (Hotelling, 1933). Two dimensions explained more than 95% of the variance in the original data. After visualizing the data in the 2 dimensions we observed 3 non-circular masses of data. Therefore, we clustered the data in 3 groups using a Gaussian Mixture Model (GMM) (McLachlan and Peel, 2000). The GMM was calculated with 3 repetitions and full covariances.

Burst rate calculation

We studied bursting by detecting short inter-spike intervals (ISIs). We detected bursts using the same procedure both at the level of single units (within or outside an ensemble) and at the level of merged spike times of all single units (within or outside an ensemble). We defined a burst as each

833 occurrence of 2 or more consecutive spikes with an ISI of less than 80 msec. We measured burst
834 rates both within periods in which the ensemble was classified as active or inactive. Burst rate was
835 measured in units of bursts per sec and was defined as the total number of bursts normalized by
836 the total time considered.

837 *LC ensemble-triggered average LFP spectral modulation*

838 We investigated the relation between the activation of LC ensembles and brain state by triggering
839 cortical LFP spectrograms on the timing of ensemble activation events. Spectra were computed
840 using the multitaper method implemented in Chronux toolbox with 3 tapers and time bandwidth
841 product of 5 (Mitra and Bokil, 2007). Short-time Fourier transforms were computed in a 10 msec
842 moving window with a duration of 200 msec. The resulting spectral resolution was ~4 Hz and the
843 temporal resolution was 10 msec. We then averaged the resulting event-triggered spectra across
844 all detected activation events separately for each ensemble. From the averaged spectra, we
845 computed an ensemble activation-triggered spectral modulation that characterized the effects of
846 LC ensemble activation on the cortical LFP power spectrum. The spectral modulation was
847 calculated as follows. We first averaged the spectrogram in time at each frequency for the baseline
848 duration (400 msec before the ensemble being active) and then subtracted the baseline averaged
849 spectrogram from the original spectrogram at each time step and divided by their sum.

850 This quantity varies between -1 to 1 for each time t and frequency f and describes the average
851 change in cortical LFP power around the time of ensemble activation.

852 *Spectrogram clustering*

853 The set of so obtained ensemble activation-triggered spectral modulations were clustered, in order
854 to assess the diversity of LC ensemble activation-triggered brain states. The clustering was
855 performed using the k-means algorithm (Arthur and Vassilvitskii, 2007). The k-means algorithm
856 requires specifying a choice for the number of possible clusters and for the mathematical function
857 used to compute the distance between the different spectrograms. We tried various definitions of
858 distance functions (Pearson Correlation, Euclidean distance, cosine, and cityblock), and we chose
859 Pearson correlation as distance function because it gave higher averaged silhouette values
860 (Rousseeuw, 1987) (i.e., cleaner clustering). We clustered the spectral modulation into k=4

clusters. This number of clusters was selected because it corresponded to the elbow point (defined as the first point in which the error drops below 5%) of the curve quantifying the normalized clustering error (error divided by the maximum error) as a function of the selected number of clusters. The error in the k-means clustering is computed as sum of the distances of each data point to their respective cluster centroid. We assessed the significance of the clustered spectral modulations at each time and frequency by pooling the spectral modulations of all ensembles in each cluster by comparing the median of the population at each point against zero using Wilcoxon signed rank test (5% significance level). The p-values were corrected for multiple comparisons using Benjamini's & Hochberg's method for false discovery rate (Benjamini and Hochberg, 1995).

The above analysis was done taking for clustering all spectral modulations obtained in correspondence of a detected activation of one or more ensembles. We performed a further control analyses in which we clustered only the subset of the spectral modulations during coactivation of ensemble-pairs. The clustering procedure for this control analysis was identical to the one reported above, but selected a number of clusters (corresponding to the elbow point of the error curve) equal to 2 clusters.

Determining the preferred cortical state for activation of different LC ensembles

We measured whether different ensembles are preferentially active in specific cortical states using state definitions from our prior work (Totah et al., 2018a). Briefly, we first divided spontaneous cortical LFPs in windows of 7.5 sec duration and then classified these windows according to their synchronization index on the basis of the relative prevalence of slow vs fast activity. Our method classified the LFP activity in each 7.5 window into one of 4 possible different categories: 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. The classification was performed as follows. 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 contained 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.

We characterized whether each LC ensemble was active during a specific cortical state by computing the likelihood of ensemble activation given a cortical state using Bayes rule. We compared this likelihood to 1000 surrogate likelihoods which were computed by shuffling the activation times (keeping the number of activation times constant). If the likelihood at a particular state crosses the 95^{th} percentile of the surrogate distribution, the ensemble was considered as preferentially 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.

References

Aghajanian GK, Cedarbaum JM, Wang RY (1977) Evidence for norepinephrine-mediated collateral inhibition of locus coeruleus neurons. Brain Res 136:570–577.

Alvarez VA, Chow CC, Bockstaele EJV, Williams JT (2002) Frequency-dependent synchrony in locus ceruleus: role of electrotonic coupling. Proc National Acad Sci 99:4032 4036.

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

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

Aston-Jones G, Bloom FE (1981b) Norepinephrine-containing locus coeruleus neurons in behaving rats exhibit pronounced responses to non-noxious environmental stimuli. J Neurosci Official J Soc Neurosci 1:887–900.

920 Aston-Jones G, Foote SL, Segal M (1985) Impulse conduction properties of noradrenergic locus
921 coeruleus axons projecting to monkey cerebrocortex. Neuroscience 15:765 777.

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

924 Aston-Jones G, Segal M, Bloom FE (1980) Brain aminergic axons exhibit marked variability in
925 conduction velocity. Brain Res 195:215–222.

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

928 Carter ME, Yizhar O, Chikahisa S, Nguyen H, Adamantidis A, Nishino S, Deisseroth K, Lecea L
929 de (2010) Tuning arousal with optogenetic modulation of locus coeruleus neurons. Nat
930 Neurosci 13:1526 1533.

931 Chandler DJ, Gao W-J, Waterhouse BD (2014) Heterogeneous organization of the locus
932 coeruleus projections to prefrontal and motor cortices. Proc National Acad Sci 111:6816
933 6821.

934 Chandler DJ, Jensen P, McCall JG, Pickering AE, Schwarz LA, Totah NK (2019) Redefining
935 Noradrenergic Neuromodulation of Behavior: Impacts of a Modular Locus Coeruleus
936 Architecture. J Neurosci Official J Soc Neurosci 39:8239–8249.

937 Chandler DJ, Lamperski CS, Waterhouse BD (2013) Identification and distribution of
938 projections from monoaminergic and cholinergic nuclei to functionally differentiated
939 subregions of prefrontal cortex. Brain Res 1522:38 58.

940 Chen F-J, Sara SJ (2007) Locus coeruleus activation by foot shock or electrical stimulation
941 inhibits amygdala neurons. Neuroscience 144:472 481.

942 Clement EA, Richard A, Thwaites M, Ailon J, Peters S, Dickson CT (2008) Cyclic and Sleep-
943 Like Spontaneous Alternations of Brain State Under Urethane Anaesthesia Greene E, ed. Plos
944 One 3:e2004.

945 Constantinople CM, Bruno RM (2011) Effects and mechanisms of wakefulness on local cortical
946 networks. Neuron 69:1061 1068.

947 Devilbiss DM, Waterhouse BD (2004) The Effects of Tonic Locus Ceruleus Output on Sensory-
948 Evoked Responses of Ventral Posterior Medial Thalamic and Barrel Field Cortical Neurons in
949 the Awake Rat. J Neurosci 24:10773–10785.

950 Ennis M, Aston-Jones G (1986) Evidence for self- and neighbor-mediated postactivation
951 inhibition of locus coeruleus neurons. Brain Res 374:299–305.

952 Eschenko O, Magri C, Panzeri S, Sara SJ (2011) Noradrenergic Neurons of the Locus Coeruleus
953 Are Phase Locked to Cortical Up-Down States during Sleep. Cereb Cortex 22:426–435.

954 Finlayson PG, Marshall KC (1988) Synchronous bursting of locus coeruleus neurons in tissue
955 culture. Neuroscience 24:217 225.

Florin-Lechner SM, Druhan JP, Aston-Jones G, Valentino RJ (1996) Enhanced norepinephrine release in prefrontal cortex with burst stimulation of the locus coeruleus. Brain Res 742:89–97.

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

Grace A, Bunney B (1984) The control of firing pattern in nigral dopamine neurons: single spike firing. J Neurosci 4:2866–2876.

Grzanna R, Molliver ME (1980) The locus coeruleus in the rat: An immunohistochemical delineation. Neuroscience 5:21 40.

Harris KD, Thiele A (2011) Cortical state and attention. Nat Rev Neurosci 12:509 523.

Hayat H, Regev N, Matosevich N, Sales A, Paredes-Rodriguez E, Krom AJ, Bergman L, Li Y, Lavigne M, Kremer EJ, Yizhar O, Pickering AE, Nir Y (2019) Locus-coeruleus norepinephrine activity gates sensory-evoked awakenings from sleep. Biorxiv:539502.

Hayat H, Regev N, Matosevich N, Sales A, Paredes-Rodriguez E, Krom AJ, Bergman L, Li Y, Lavigne M, Kremer EJ, Yizhar O, Pickering AE, Nir Y (2020) Locus coeruleus norepinephrine activity mediates sensory-evoked awakenings from sleep. Sci Adv 6:eaaz4232.

Hotelling H (1933) Analysis of a complex of statistical variables into principal components. J Educ Psychol 24:417–441.

Howe M, Ridouh I, Mascaro ALA, Larios A, Sedano MA, Dombeck D (2019) Coordination of rapid cholinergic and dopaminergic signaling in striatum during spontaneous movement. Elife 8:e44903.

Huang H-P, Wang S-R, Yao W, Zhang C, Zhou Y, Chen X-W, Zhang B, Xiong W, Wang L-Y, Zheng L-H, Landry M, Hökfelt T, Xu Z-QD, Zhou Z (2007) Long latency of evoked quantal transmitter release from somata of locus coeruleus neurons in rat pontine slices. Proc National Acad Sci 104:1401 1406.

Ishimatsu M, Williams JT (1996) Synchronous Activity in Locus Coeruleus Results from Dendritic Interactions in Pericoerulear Regions. J Neurosci 16:5196–5204.

Lee A, Rosin DL, Bockstaele EJV (1998) α2A-adrenergic receptors in the rat nucleus locus coeruleus: subcellular localization in catecholaminergic dendrites, astrocytes, and presynaptic axon terminals. Brain Res 795:157–169.

Lee DD, Seung HS (1999) Learning the parts of objects by non-negative matrix factorization. Nature 401:788–791.

Lovett-Barron M, Andalman AS, Allen WE, Vesuna S, Kauvar I, Burns VM, Deisseroth K (2017) Ancestral Circuits for the Coordinated Modulation of Brain State. Cell 171:1411 1423.e17.

992 Martins ARO, Froemke RC (2015) Coordinated forms of noradrenergic plasticity in the locus
993 coeruleus and primary auditory cortex. Nat Neurosci 18:1483 1492.

994 Marzo A, Totah NK, Neves RM, Logothetis NK, Eschenko O (2014) Unilateral electrical
995 stimulation of rat locus coeruleus elicits bilateral response of norepinephrine neurons and
996 sustained activation of medial prefrontal cortex. J Neurophysiol 111:2570–2588.

997 McCormick DA (1992) Neurotransmitter actions in the thalamus and cerebral cortex and their
998 role in neuromodulation of thalamocortical activity. Prog Neurobiol 39:337–388.

999 McCormick DA, Nestvogel DB, He BJ (2020) Neuromodulation of Brain State and Behavior.
1000 Annu Rev Neurosci 43:1–25.

1001 McCune SK, Voigt MM, Hill‡ JM (1993) Expression of multiple alpha adrenergic receptor
1002 subtype messenger RNAs in the adult rat brain. Neuroscience 57:143–151.

1003 McGinley MJ, Vinck M, Reimer J, Batista-Brito R, Zagha E, Cadwell CR, Tolias AS, Cardin JA,
1004 McCormick DA (2015) Waking State: Rapid Variations Modulate Neural and Behavioral
1005 Responses. Neuron 87:1143 1161.

1006 McLachlan G, Peel D (2000) Wiley Series in Probability and Statistics.

1007 Mitra P, Bokil H (2007) Observed Brain Dynamics. Oxford University Press.

1008 Navarra RL, Clark BD, Zitnik GA, Waterhouse BD (2013) Methylphenidate and atomoxetine
1009 enhance sensory-evoked neuronal activity in the visual thalamus of male rats. Exp Clin
1010 Psychopharm 21:363 374.

1011 Neves RM, Keulen S van, Yang M, Logothetis NK, Eschenko O (2018) Locus Coeruleus phasic
1012 discharge is essential for stimulus-induced gamma oscillations in the prefrontal cortex. J
1013 Neurophysiol 119:jn.00552.2017.

1014 Onken A, Liu JK, Karunasekara PPCR, Delis I, Gollisch T, Panzeri S (2016) Using Matrix and
1015 Tensor Factorizations for the Single-Trial Analysis of Population Spike Trains. Bethge M, ed.
1016 Plos Comput Biol 12:e1005189.

1017 Paxinos G, Watson C (2017) The Rat Brain in Stereotaxic Coordinates: Compact 7th Edition.
1018 Academic Press.

1019 Polack P-O, Friedman J, Golshani P (2013) Cellular mechanisms of brain state-dependent gain
1020 modulation in visual cortex. Nat Neurosci 16:1331 1339.

1021 Rajkowski J, Kubiak P, Aston-Jones G (1994) Locus coeruleus activity in monkey: Phasic and
1022 tonic changes are associated with altered vigilance. Brain Res Bull 35:607 616.

1023 Rajkowski J, Majczynski H, Clayton E, Aston-Jones G (2004) Activation of monkey locus
1024 coeruleus neurons varies with difficulty and performance in a target detection task. J
1025 Neurophysiol 92:361 371.

Rand WM (1971) Objective Criteria for the Evaluation of Clustering Methods. J Am Stat Assoc 66:846.

Rousseeuw PJ (1987) Silhouettes: A graphical aid to the interpretation and validation of cluster analysis. J Comput Appl Math 20:53–65.

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

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

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

Swanson LW (1976) The locus coeruleus: a cytoarchitectonic, Golgi and immunohistochemical study in the albino rat. Brain Research 110:39 56.

Swanson LW, Hartman BK (1975) 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.

Totah NK, Logothetis NK, Eschenko O (2015) Atomoxetine accelerates attentional set shifting without affecting learning rate in the rat. Psychopharmacology 232:3697–3707.

Totah NK, Neves RM, Panzeri S, Logothetis NK, Eschenko O (2018a) The Locus Coeruleus Is a Complex and Differentiated Neuromodulatory System. Neuron 99:1055-1068.e6.

Totah NKB, Logothetis NK, Eschenko O (2018b) Noradrenergic ensemble-based modulation of cognition over multiple timescales. Brain Res.

Tung C-S, Ugedo L, Grenhoff J, Engberg G, Svensson TH (1989) Peripheral induction of burst firing in locus coeruleus neurons by nicotine mediated via excitatory amino acids. Synapse 4:313–318.

Waterhouse BD, Woodward DJ (1980) Interaction of norepinephrine with cerebrocortical activity evoked by stimulation of somatosensory afferent pathways in the rat☆. Exp Neurol 67:11–34.

Williams AH, Kim TH, Wang F, Vyas S, Ryu SI, Shenoy KV, Schnitzer M, Kolda TG, Ganguli S (2018) Unsupervised Discovery of Demixed, Low-Dimensional Neural Dynamics across Multiple Timescales through Tensor Component Analysis. Neuron 98:1099-1115.e8.

Acknowledgements

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

Author Contributions

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

Declaration of Interests

The authors declare no competing interests.

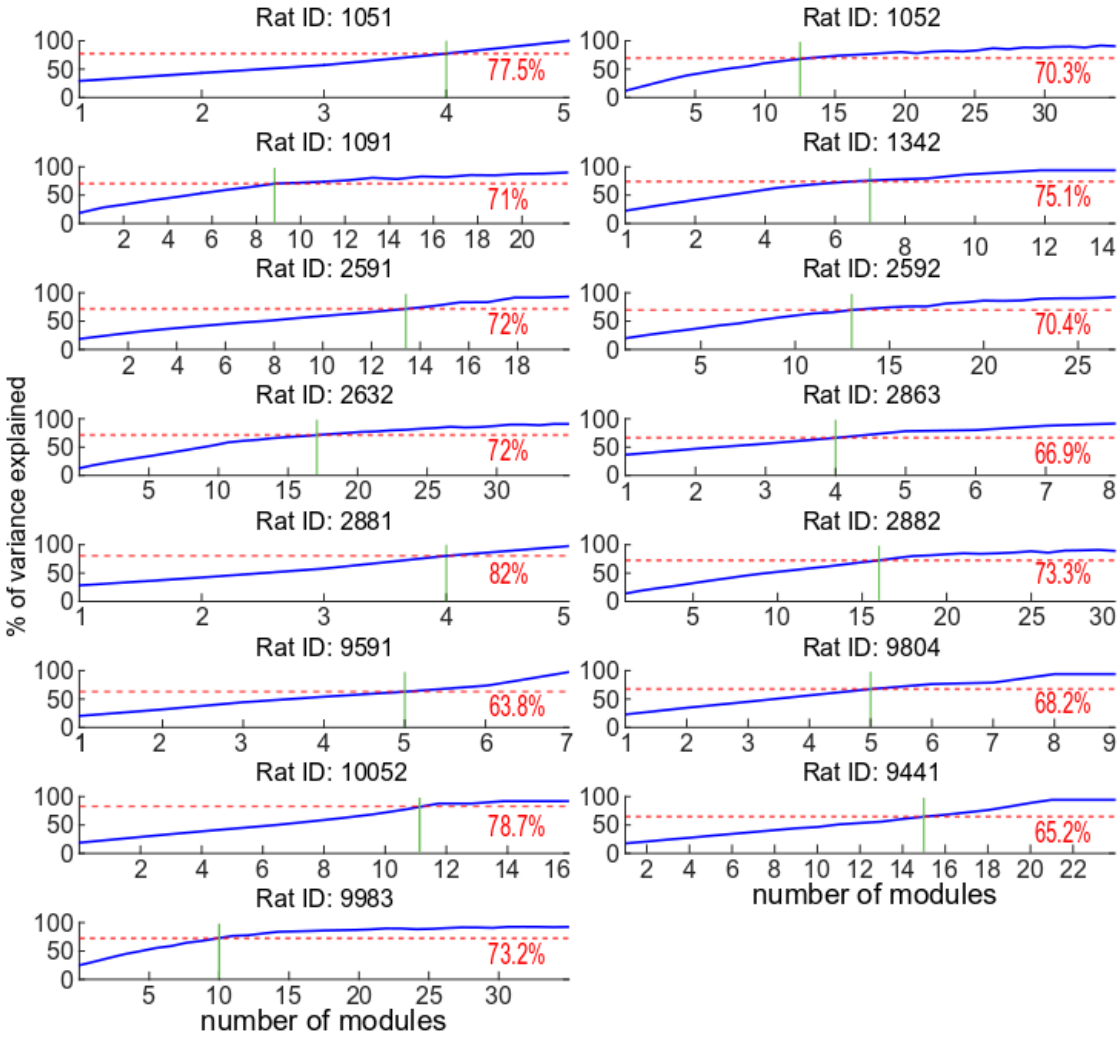

Supplementary Figure 1. Data underlying the choice of the optimal number of modules (K) in each rat. Each panel depicts the percentage of explained variance versus the number of the modules for each rat. Solid green lines show the number of selected modules based on the criteria of first elbow after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

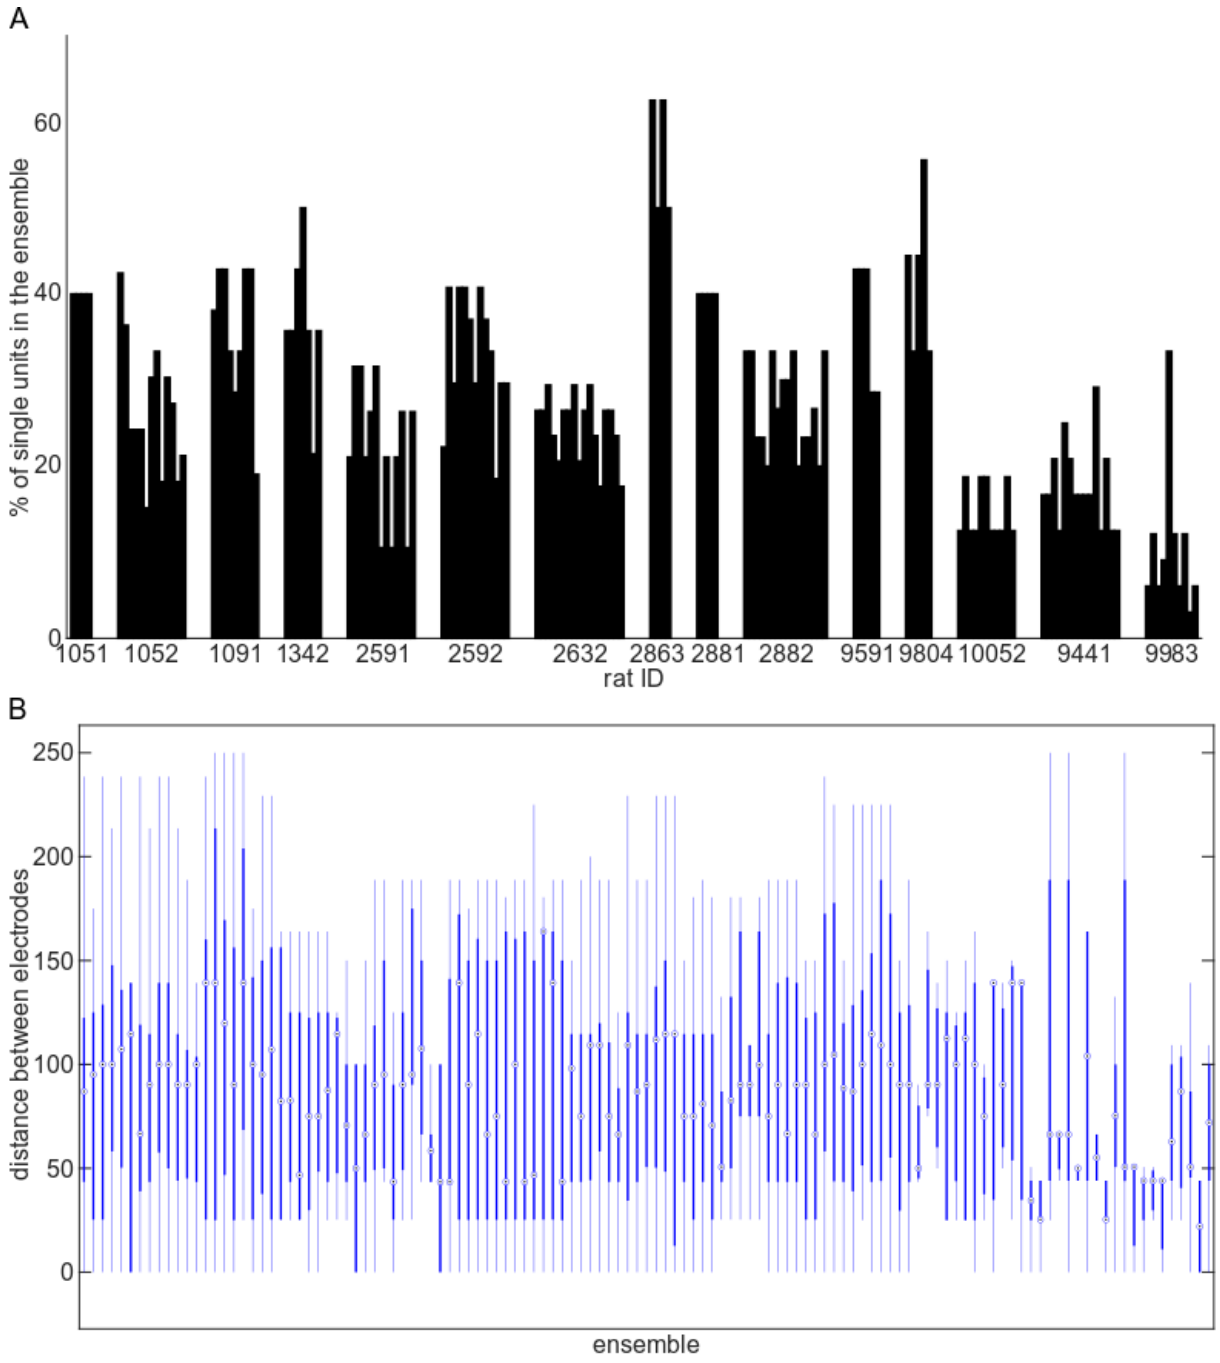

1079

1080 **Supplementary Figure 2. The spatial properties of the detected ensembles. (A)** Bar plot
1081 showing the percentage of all simultaneously recorded single units within each ensemble. The
1082 percentage was calculated as the number of single units inside the ensemble divided by the total
1083 number of single units recorded for that rat. Each ensemble is a bar. The bars are grouped by rat.
1084 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.

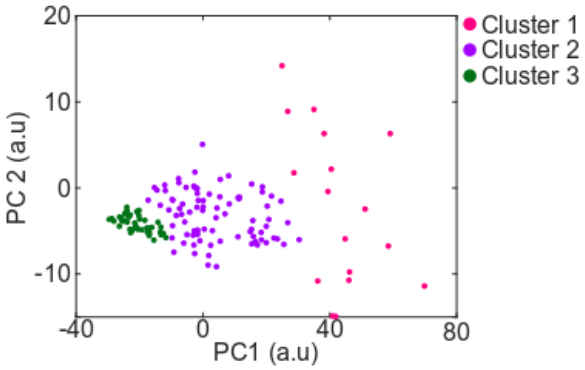

Supplementary Figure 3. Supporting data showing the clustering of PETHs. The scatter plot shows the projections of the PETHs into two dimensions (PC1, PC2) using PCA. The first two principle components explained more than 95% of the variance. Three non-circular masses were clustered using GMM. Data points falling into each cluster are color coded separately.

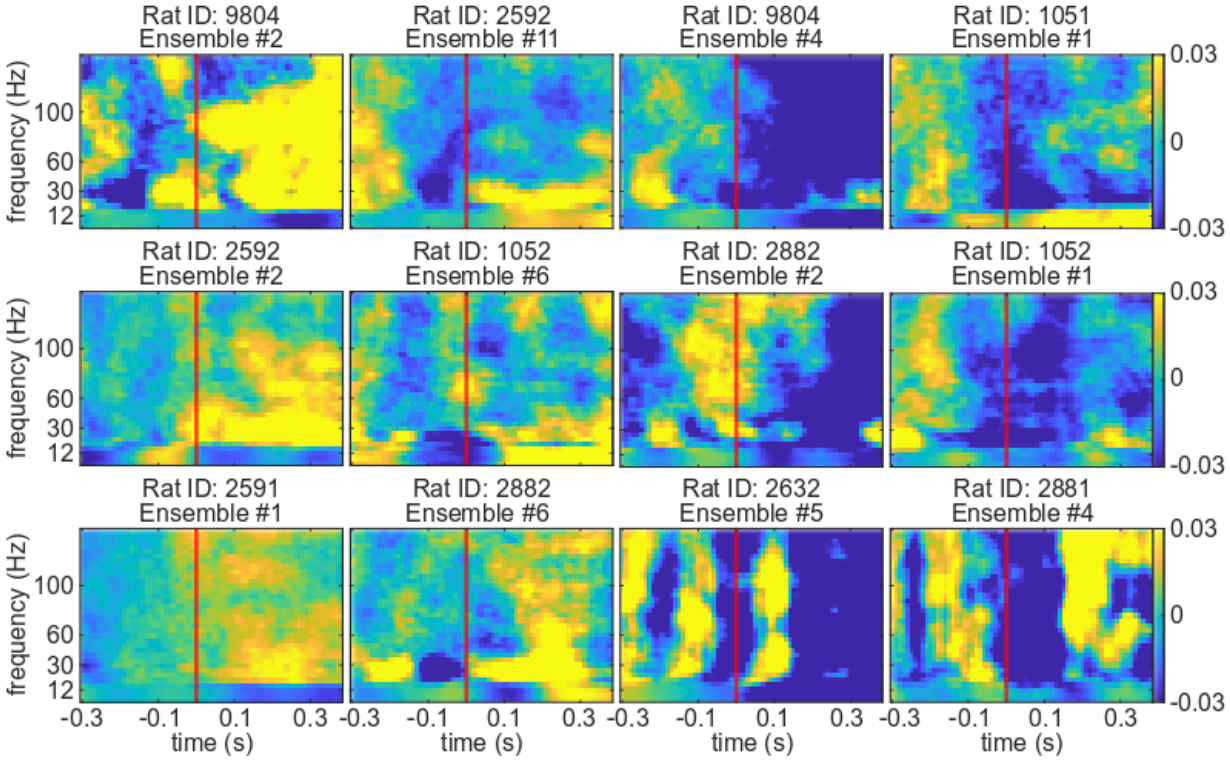

1095

1096 **Supplementary Figure 4. Examples of LC ensemble activation-triggered cortical LFP power**
1097 **spectra.** Examples from 12 different LC ensembles illustrate the diverse cortical states which
1098 occur around the time of ensemble activation. The examples are shown in 4 columns, each of
1099 which indicates a specific trend in the spectra corresponding to the clusters shown in the Figure 5.

1100

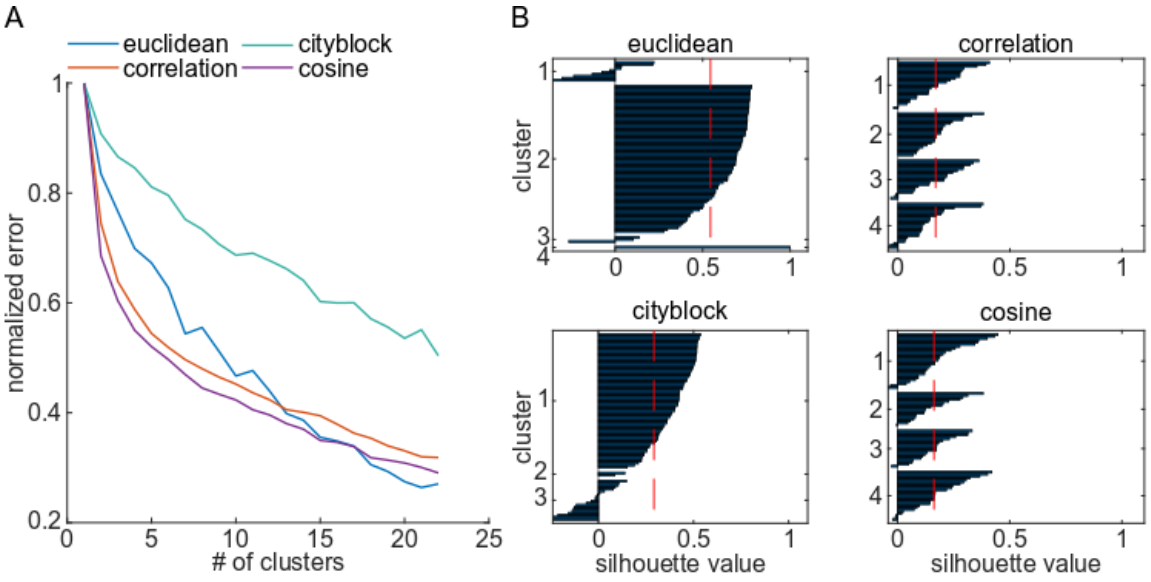

Supplementary Figure 5. The result of analyses supporting the determination of the best criteria for spectral clustering. (A) The normalized error (error divided by the maximum error) of the k-means clustering of the ensemble activation-triggered spectra versus the number of clusters. Four different 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).