Bridging scales from spiking activity to the local field potential through latent dynamics

Audrey Sederberg[1,2], Aurélie Pala[1], Garrett B Stanley[1*]

[1]Wallace H Coulter Department of Biomedical Engineering, Georgia Institute of Technology and Emory University, Atlanta, GA, USA, [2]Department of Neuroscience, University of Minnesota, Minneapolis, MN

Running headline: Multiscale dynamics of cortical states
With 47 text pages, 8 figures

*Correspondence: Audrey Sederberg

                   Department of Neuroscience

                   University of Minnesota

                   2021 6th Street SE

                   Minneapolis, MN 55455

                   Phone: 612-626-3891

                   Email: sede0018@umn.edu

Keywords: **Cortical state; LFP; population spiking; latent dynamics; variability**

Funding: This work was supported by the National Institutes of Health (NIH BRAIN Grants R01NS104928, U01NS094302, R01EB029857 to GBS, NIH R21NS112783 to GBS and AP); Swiss National Science Foundation postdoctoral fellowships P2ELP3_168506 and P300PA_177861 (AP). Additional support was provided to AS through startup funds from the University of Minnesota Medical School.

Abstract (219 words)

During wakefulness, cortical networks exhibit diverse states of activity arising from fluctuations in modulatory inputs, any of which may shape the recorded neural activity across a range of scales, from local field potential (LFP) to the spiking of individual neurons. Multiple studies have uncovered state-dependent changes in single neuron activity specific to cell type. These changes are heterogeneous across cell types, suggesting that there is a link between changes in states of wakefulness and the dynamic patterns of spiking activity across the population, yet this is currently not well understood. Here, we use an unsupervised statistical modeling framework to link fluctuations in the state of the LFP to the dynamics of spiking activity of populations of neurons in the whisker area of primary somatosensory cortex of awake male mice. By identifying latent states from population spiking activity, we show that these spiking states are more informative of states of the LFP than a direct readout of the spiking activity of single neurons or of the population. Further, we find that the latent spiking states predict trial-by-trial variability in sensory responses as well as one aspect of behavior, whisking activity. Taken together, our approach captures the latent variables that underlie the dynamics of the spiking activity of cortical neuronal populations, providing a framework to relate cortical activity across different scales.

Significance statement (118 words)

During wakefulness, variability in cortical activity can arise from shifts in arousal or behavior. We analyze the dynamics of two measures of cortical activity, the spiking activity of local populations of neurons and the local field potential (LFP), typically used to track overall levels of arousal. We show that the latent states that capture variability in population spiking activity are highly predictive of LFP state changes, even more so than spiking activity directly. By

linking latent states of spiking activity to changes in the LFP, our novel approach bridges fine-scale microcircuits ensemble activity and the meso-scale LFP. Our latent-state framework opens new avenues for extracting information from the LFP as well as identification of subject-level sources of variability.

Introduction (648 words)

Brain activity is dynamic and variable, even in the absence of sensory drive or task demands, arising spontaneously from internal sources. A large body of literature relates fluctuations in metrics of cortical activity to fluctuations in cognitive and sensorimotor behaviors, in humans (Monto *et al.*, 2008; He and Zempel, 2013) as well as in animal models (Pinto *et al.*, 2013; Engel *et al.*, 2016; Beaman, Eagleman and Dragoi, 2017; Speed *et al.*, 2019; McCormick, Nestvogel and He, 2020; Salkoff *et al.*, 2020). Due to the close relationship between ongoing cortical activity and perception and cognition, understanding the multi-scale organization of cortical dynamics is expected to provide an invaluable predictor of how individual variability in cortical function arises from differences in microcircuit activity. Making this advance requires data-driven, quantitative models linking the dynamics of an ensemble of neurons in a local microcircuit to meso-scale measures, such as the local field potential (LFP).

Bridging from spiking activity to the LFP is not simple. The LFP is a collective readout of synaptic activity in a population of neurons, and in principle, detailed biophysical microcircuit models can predict the LFP from layer- and cell-type-specific population spiking activity (Gratiy *et al.*, 2011; Hagen *et al.*, 2016). However, this problem is severely under-constrained: there are too many unknowns in any microcircuit to directly predict population activity. A potential resolution comes from modeling ensemble dynamics in the awake brain, which has advanced with the development of experimental methods to monitor large populations of neurons

simultaneously. Modeling studies have suggested that a relatively small number of latent dynamical variables control much of the moment-by-moment variability evident in single-neuron spiking activity (Köster *et al.*, 2014; Okun *et al.*, 2015). The importance of these variables in predicting population activity is apparent in the growing set of studies uncovering statistical links between behavior and large-scale population activity in the mouse (Musall *et al.*, 2019; Stringer *et al.*, 2019). Thus, identifying latent variables in ensemble recordings may better constrain efforts to link spiking activity to the LFP.

Neural activity during wakefulness in the mouse cortex has been extensively characterized using the LFP and the activity patterns of single neurons of identified cell types (Petersen *et al.*, 2003; Poulet and Petersen, 2008; Bennett, Arroyo and Hestrin, 2013; Eggermann *et al.*, 2014; Vinck *et al.*, 2015; Fernandez *et al.*, 2016; Pala and Petersen, 2018; Sederberg *et al.*, 2019; Senzai, Fernandez-Ruiz and Buzsáki, 2019). For instance, in the somatosensory cortex, "state" changes in spontaneous activity may be described on a continuum, from "deactivated" to "activated" (Poulet and Crochet, 2019), with the deactivated state characterized by local field potential (LFP) activity dominated by low-frequency components (1-10 Hz). Pre-stimulus LFP state is predictive of trial-by-trial variability in the sensory response (Haslinger *et al.*, 2006; Sederberg *et al.*, 2019), and LFP state is informative of behavioral metrics, including whisking activity (Fernandez *et al.*, 2016). Studies of single neurons revealed cell-type-specific changes in firing rates with fluctuations in activation state (Gentet *et al.*, 2010, 2012; Pala and Petersen, 2018). These results suggest that the collective population dynamics are a rich source of information about ongoing LFP state, and vice versa.

Here, we first adapt a standard approach for analysis of LFP state (Fernandez *et al.*, 2016; Poulet and Crochet, 2019) to account for variability across individuals. We then fit spontaneous

population spiking activity with a model of latent dynamics, extracting a "latent state" representation of population activity. We examine the relationship between latent spiking state dynamics and the LFP state, finding distinct sets of spiking states depending on the LFP states. Finally, we analyze the organization of latent states, showing that they predict trial-by-trial variability in sensory-evoked responses as well as the intensity of ongoing whisking activity. Our results demonstrate how modeling latent dynamics provides a systematic, quantitative framework to characterize cortical dynamics, important for future investigations of cortical activity dynamics across diverse behavioral and cognitive states.

Materials and Methods

Experimental data collection

Data presented here were previously reported in (Pala and Stanley, 2022) and (Sederberg *et al.*, 2019). All procedures were approved by the Institutional Animal Care and Use Committee at the Georgia Institute of Technology and agreed with guidelines established by the National Institutes of Health. Four nine- to twenty-six-week-old male C57BL/6J mice were used in this study. Mice were maintained under 1-2% isoflurane anesthesia while being implanted with a custom-made head-holder and a recording chamber. The location of the barrel column targeted for recording was functionally identified through intrinsic signal optical imaging (ISOI) under 1-1.25% isoflurane anesthesia. Recordings were targeted to B1, B2, C1, C2, and D2 barrel columns. Mice were habituated to head fixation, paw restraint and whisker stimulation for 3-6 days before proceeding to electrophysiological recordings.

Electrophysiology

Local field potential and spiking activity were recorded using silicon probes (A1x32-5mm-25-177, NeuroNexus, USA) with 32 recording sites along a single shank covering 775 µm in depth. The probe was coated with DiI (Invitrogen, USA) for post hoc identification of the recording site. The probe contacts were coated with a poly(3,4-160 ethylenedioxythiophene) (PEDOT) polymer (Ludwig *et al.*, 2006) to increase signal-to-noise ratio. Contact impedance measured between 0.3 MOhm and 0.5 MOhm. The probe was inserted with a 35° angle relative to the vertical, until a depth of about 1000 µm. Continuous signals were acquired using a Cerebus acquisition system (Blackrock Microsystems, USA). Signals were amplified, filtered between 0.3 Hz and 7.5 kHz and digitized at 30 kHz.

Recording of spontaneous and sensory-evoked activity

Mechanical stimulation was delivered to a single contralateral whisker corresponding to the barrel column identified through ISOI using a galvo motor (Cambridge Technologies, USA).

The galvo motor was controlled with millisecond precision using a custom software written in Matlab and Simulink Real-Time (Mathworks, USA). The whisker stimulus followed a sawtooth waveform (16 ms duration) of various velocities (1000 deg/s, 500 deg/s, 250 deg/s, 100 deg/s) delivered in the caudo-rostral direction (Wang, Webber and Stanley, 2010). Whisker stimuli of different velocities were randomly presented in blocks of 21 stimuli, with a pseudo-random inter-stimulus interval of 2 to 3 seconds and an inter-block interval of a minimum of 20 seconds. The total number of trials across all velocities presented during a recording session ranged from 196 to 616. Recordings were acquired in 10-minute segments, and the end of each 10-minute segment was 1 minute of spontaneous activity. All spontaneous segments of 10 seconds or longer were analyzed.

LFP pre-processing

For analysis, the LFP was down-sampled to 2 kHz. Using the same quality criterion as previously (Sederberg *et al.*, 2019), we screened for channels that drifted or had excessive 60-Hz line noise. Only layer-4 LFP was analyzed in this study. Layer 4 was determined using the sensory-evoked MUA activity and current source density profiles as described previously (Sederberg *et al.*, 2019; Pala and Stanley, 2022).

Spike sorting and single unit quality criteria.

Spike sorting was performed using Kilosort2 (Pachitariu *et al.*, 2016) with manual curation of spike clusters in Phy (Rossant and Harris, 2013). The quality criteria for identifying single unit activity were the same as in (Pala and Stanley, 2022), summarized by these six requirements: (i) units had more than 500 spikes, (ii) spike waveform SNR was greater than 5, (iii) the trough-to-peak amplitude of the unit had a coefficient of variation less than 0.2 averaging over 120-s windows, (iv) firing rate (over 120-s windows) had a coefficient of variation of less than 1, (v) fewer that 1% of inter-spike intervals were shorter than 2 ms, and (vi) the Maholonobis

distance between the unit waveform and the closest waveform that was *not* assigned to that unit was larger than 55 in the space of the first three principal components.

Units were classified as RS or FS based on the time elapsed from the waveform trough to waveform peak, with values larger than 0.5 ms classified as RS units and those smaller than 0.4 as FS units (Barthó *et al.*, 2004). Intermediate waveforms (about 3% of units) were not included in the analysis. Overall, we analyzed 39 FS units and 119 RS units across six recordings (average of 26 units per recording, range 19 to 36) in four different mice.

Whisking quantification.

Video recordings of whisking activity were acquired under infrared illumination at 25 Hz using a camera (6.8 pixels/mm, HQCAM) positioned below the mouse's head, acquiring a view of whiskers as well as of the galvanometers used for stimulus control. For each movie, the portion of the field of view including the whiskers was extracted. Whisking activity was extracted by first identifying motion components using non-negative matrix factorization of the square difference between frames. Up to 20 components were identified and these were manually scored as representing whisker motion or galvanometer motion. Galvo traces were used to verify temporal alignment of movies and electrophysiology. Whisking activity was quantified as the sum over whisker-motion components.

Experimental design and statistical analysis

Multiple statistical analyses were used in this study and are described in detail below. All statistical analyses were conducted in MATLAB (Mathworks).

LFP and Spiking State Model Fitting and Validation

Fitting the mixture of Gaussians model for LFP.

We calculated LF/HF, following a convention used in several earlier studies (Fernandez *et al.*, 2016; Poulet and Crochet, 2019). Specifically, we compute the spectrogram at 1-s intervals

and defining LF (HF) as the average Fourier amplitude between 1 and 10 Hz (30 and 90 Hz, excluding 58 to 62 Hz). We consider the distribution of the transformed variable

$$r = -1 + \log_{10} \frac{\text{LF}}{\text{HF}}$$

Thus, $r > 0$ corresponds to LF/HF > 10 and $r < 0$ corresponds to LF/HF < 10. and fit the histogram of r for each recording to a mixture of Gaussians model

$$p(r) = \sum_{i=1}^{N_g} \frac{1}{\sqrt{2\pi\sigma_i^2}} \exp \frac{(r - \mu_i)^2}{2\sigma_i^2}$$

We calculated the Bayesian Information Criterion (BIC) to determine the number of components.

Hidden Markov Model of cortical spiking activity

To describe cortical spiking states, we fit a Hidden Markov Model with Poisson emission distribution (Radons, Becket and Diilfer, 1994; Abeles *et al.*, 1995; Jones *et al.*, 2007; Mazzucato, Fontanini and La Camera, 2015). This model assumes that the current state $s_t \in [1,2,...N_S]$ only depends on the previous state s_{t-1}, and this is characterized by the transition matrix A:

$$A_{ij} = p(s_t = j | s_{t-1} = i)$$

The emission probabilities capture the distribution of spike counts in a time bin of size Δt; we assuming firing is Poisson and conditionally independent across neurons. Specifically, the probability of observing a spike count $y_i(t) = k$ for neuron i in state $s_t = m$ between time t and $t + \Delta t$ is

$$P(y_i = k | s_t = m) = \frac{(\lambda_i(m)\Delta t)^k e^{-\lambda_i(m)\Delta t}}{k!} \qquad (1)$$

The joint probability across a set of C simultaneously recorded cells is

$$P(\boldsymbol{y} = \boldsymbol{k}|s_t = m) = \prod_{i=1}^{C} P(y_i = k_i|s_t = m)$$

Thus, each state is defined by a vector of firing rates $\lambda(m)$. We binned spike counts at a temporal resolution of 40 ms to match the resolution (25 Hz) of video monitoring of whisking activity. In other works, the temporal resolution of spiking states has ranged from 1 ms (Jones *et al.*, 2007) to 500 ms (Abeles *et al.*, 1995).

The model was implemented by adding a conditionally independent Poisson observation model to the pmtk-3 (probabilistic models toolbox, https://github.com/probml/pmtk3). To fit the model, we split each recording into two halves as follows. Each recording consisted of a set of 8 to 12 segments, each of 10 minutes duration. Each of these long segments had 10 periods of spontaneous activity (duration, 10 s to 70 s) with no controlled whisker stimulation. Model parameters were estimated on a training set consisting of spontaneous periods from half of the long segments (1st, 3rd, 5th, etc.). Using the estimated model for each recording, states $s_i(t)$ at time t in each spontaneous period in the test set (the remaining half of long segments: 2nd, 4th, 6th, etc.) were decoded by estimating $p(s_i(t)|\{y(1:T))$ using standard methods (forward-backward algorithm). Reported likelihood values are computed across all spontaneous periods in the test set.

We used two controls to interpret the data likelihood under the fitted model. First, to estimate how well the model could have done had the true generating process matched the fitted model, we simulated an HMM using the transition matrix fitted, and drew from the emission distribution at each point in time according to the simulated state. We then computed the likelihood of this simulated data under the fitted model (Fig. 4). Second, we verified that we could not obtain the same result in shuffled data. Thus, we shuffled the simulated observations in time and recomputed the likelihood of shuffled observations under the fitted model.

For the analysis of pre-stimulus state, we extracted the sensory-evoked periods of each long recording, and using the 500-ms pre-stimulus time window, we inferred HMM state using the model previously fit on the spontaneous periods. The pre-stimulus state was assigned to the state with the highest probability in the final pre-stimulus time bin (-40 ms to 0 ms pre-stimulus).

Quantifying differences between spiking states

To analyze the differences between states identified by the spiking state model, Bhattacharyya coefficients and distances were computed for pairs of spiking state emission distributions. The Bhattacharyya coefficient between two (potentially non-Gaussian) distributions quantifies their overlap and determines the Bhattacharyya distance D_{BC} through the relationship $BC = \exp(-D_{BC})$ (Bhattacharyya, 1946). For each cell i and pair of states (a, b), the state-dependent emission distributions are Poisson distributions of rates $\lambda_i(a)$ and $\lambda_i(b)$ (as in Equation (1)), for which the Bhattacharyya distance is

$$D_{BC}(a, b\,; i) = \frac{1}{2}\left(\sqrt{\lambda_i(a)} - \sqrt{\lambda_i(b)}\right)^2 \Delta t$$

and Δt is the time bin (40 ms). Because we have assumed conditional independence within each state, the full distribution across cells is a product distribution, and the distances add across cells. The Bhattacharyya coefficient between states a and b is then

$$BC(a, b) = \exp\left(-\sum_i D_{BC}(a, b; i)\right)$$

This ranges from 0 to 1, where 1 indicates complete overlap $(\lambda_i(a) = \lambda_i(b))$ and 0 is no overlap.

Analyses of inferred spiking states

Ordering spiking states by LFP states

We ordered spiking states based on whether LFP state H or LFP state L was more likely to be observed during the periods in which that spiking state was inferred (Fig. 5). Specifically, we computed $P(x|s_i)$, where s_i is the spiking state and x is the LFP state (either H, if LF/HF < 10, or else L). The "LFP State ratio" (r_{LS},) is defined for each state s_i as

$$r_{LS} = \frac{P(H|s_i)}{P(L|s_i)}.$$

Decoding analysis of LFP-spiking state relationships

Multiple classifiers were fit in Fig. 6. For all classifiers, the target response is whether LF/HF > 10. Observations are balanced so that 50% accuracy is chance. The inputs to the different types of classifiers are (i) summed spike count across the population, (ii) the population spike count vector, (iii) the inferred spiking state, or (iv) the single unit spike counts (Fig. 6B). Observations were divided into test and training sets. The classifier was a linear SVM.

Spiking state-dependent analysis of sensory-evoked responses

The evoked response is defined as the difference between the spike count in the post-stimulus window $[5\text{ms}, 25\text{ms}]$ and the pre-stimulus window $[-20\text{ms}, 0\text{ms}]$. Evoked spike counts were computed for each trial, and we wanted to know if the average evoked response in each state was higher or lower than would be expected from the null hypothesis that evoked responses have no dependence on pre-stimulus spiking state. The number of trials assigned to each state was not identical, making direct comparison difficult. We determined significance using a shuffle test. Specifically, we randomly assigned trials to states matching the number of trials that were classified to that state and computed the average evoked rate by state in each random draw. This was repeated 200 times for each recording and for each number of spiking states (2 to 10) to generate the shuffled distribution of average evoked response. The

empirical mean state-dependent evoked spike count was compared to this distribution. Significance was determined by a two-tailed 1% cut-off, so 'facilitated' responses were those in which the empirical mean state-dependent evoked spike count scored above the 99.5% percentile of all shuffled mean evoked responses.

Finally, we assessed whether the state-dependent single unit responses were organized across the population or randomly distributed. We therefore tested whether the number of single units with significantly different evoked responses tended to be clustered in a single state. For each single unit i and each state s, we label the result σ_{is} of this test by a $+1$ (evoked response was higher than expected), a -1 (evoked response was lower than expected), or a 0 (evoked response was within the shuffle distribution). We first address whether we observe more $+1$'s in a single state than would be expected by chance, but the calculation is similar for the -1 case. Thus, we define n_i^+ to be the number of $+1$'s in the state with the most $+1$'s and N^+ to be the number of single units that have a $+1$ in *any* state. If every single unit had the same number of $+1$'s, then we would simply count the number of combinations $\binom{N^+}{k}$ of all single units with k entries equal to $+1$ in the state of interest, summing from $k = n_i^+ \dots N^+$. Empirically, not all units have the same number of $+1$'s, so we compute p_i^+ as the fraction of states for which single unit i had a $+1$ and include in our count a weighting factor for the different probabilities. Specifically, we compute the probably of a state having n_i^+ or more single units with a $+1$ by summing over all combinations $\tilde{\sigma}_{is}^j$ with k total $+1$'s in state s:

$$P(k \geq n_i^+; p_i^+) = \sum_{k=n_i^+}^{N^+} \sum_{\{\sum_i \tilde{\sigma}_{is}^j = k\}} \prod_i (p_i^+)^{\tilde{\sigma}_{is}^j} (1 - p_i^+)^{1 - \tilde{\sigma}_{is}^j}.$$

Here $\tilde{\sigma}_{is}^j \in \{0,1\}$ indicates whether the entry is a 1 (i.e., $\tilde{\sigma}_{is}^j = 1$) or not a 1. This calculation is repeated for -1, in which case $\tilde{\sigma}_{is}^j = 1$ indicates that the i, s entry is a -1.

Information calculation between whisking activity and spiking state identity

Mutual information quantifies the difference between the joint distribution $P(X, Y)$ of quantities X and Y compared to the distribution $P(X)P(Y)$ in which X and Y are independent, providing a generalized measure of correlation (including non-linear correlation) between quantities X and Y. Specifically, the mutual information (in bits) is $\langle \log_2 \frac{P(X,Y)}{P(X)P(Y)} \rangle$, where angle brackets denote an average over the joint distribution $P(X, Y)$. For each calculation of mutual information (Fig. 8), X is the whisking energy, computed in each 40-ms time bin (i) from videography and discretized into four equipartition bins ($W_i \in \{0,1,2,3\}$). To compute information between whisking and spiking states ("Y" is the spiking state), we identified the maximum-likelihood spiking state ($s_{spk}(i) \in \{0, 1, \dots N_S\}$) in each time bin and estimated the mutual information between W_i and $s_{spk}(i)$. To compute information between whisking energy and LFP state ("Y" is LFP state), we assigned LFP state to L or H according to whether $r>0$ or $r<0$. Error bars were estimated by resampling over half the data.

Code accessibility

Custom MATLAB code to generate figures and perform analyses can be accessed from GitHub following publication.

Results

We developed a data-driven modeling framework to track the dynamics of cortical state, at the coarse-grained scale of the local field potential (LFP) as well as the fine-grained scale of spiking activity in simultaneously recorded single neurons. We acquired recordings of the LFP and single-unit spiking activity using linear multi-electrode arrays (Neuronexus A1x32-5mm-25-177-A32) in the primary somatosensory cortex (S1) of the awake, head-fixed mouse (Fig. 1A) while monitoring whisking activity with videography (Methods).

As expected from previous studies (Fernandez *et al.*, 2016; Poulet and Crochet, 2019), during a typical epoch of spontaneous LFP activity in the absence of experimenter-controlled sensory input, the local cortical network spontaneously transitioned between periods of stronger high-frequency activity (Fig. 1B, exemplar epoch marked by black bar) and periods of stronger low-frequency activity (Fig. 1B, exemplar epoch marked by green bar). Our goal was to first assess how well the standard approach of using the ratio of low- to high-frequency LFP power for classifying state transitions using LFP in the awake mouse captured state transitions over the course of a recording, and then to relate signatures of LFP state to collective population activity patterns through a hidden-state model of spiking activity. Using this framework, we examined the extent to which the underlying spiking states predict trial-by-trial sensory-evoked variability and whisking activity.

Figure 1. LFP states during wakefulness.
A: Local field potential (LFP) and single unit (SU) activity were recorded from laminar array electrodes in primary somatosensory (S1) cortex of awake head-fixed mouse.
B: LFP in layer 4 over a 20-s period. Time axes in B and D are the same. Black bar: example of *activated* state ("H"). Green bar: example of the *deactivated* state ("L").
C: Fourier amplitude of the LFP. LF is the average amplitude between 1 and 10 Hz and HF is the average amplitude between 30 and 90 Hz. Fourier amplitude spectrum averaged over activated (black) and inactivated (green) periods, where states were classified based on LF/HF ratio.
D: LF/HF evaluated in sliding 1-s window over 20 seconds, with black/green shading corresponding to bars in (B).
E: Density plot (color represents log density) of the cortical state metric ($r = log (LF/HF) - 1$) and whisker motion, quantified from video recording (see Methods). Linear correlation was -0.43 (p = 1e-22).
F: Distribution of r and fit to a mixture of Gaussians model with two components.
G: Bayesian Information Criterion (BIC) for mixture of Gaussians model for 1 to 8 states is minimized at n = 2 states.

Statistics of LFP metrics of cortical state in S1 of awake mice

First following the choice of metric used in past studies of mouse S1 (Fernandez *et al.*, 2016; Poulet and Crochet, 2019; Sederberg *et al.*, 2019), the observable quantity driving the analysis was based on the power spectral density derived from the Fourier transform of the LFP (Fig. 1C). The power spectrum was integrated over two frequency bands and denoted by LF (power in the 'low frequency' band, 1-10 Hz) and HF (power in the 'high frequency' band, 30 – 90 Hz, excluding 58 to 62 Hz). In previous studies, a heuristic for classifying epochs as strong HF

("activated") or strong LF ("deactivated") was suggested based on the ratio of the power in the two bands, denoted here as LF/HF. Thus, one can classify epochs as "activated" (LF/HF < 10, black curve, Fig. 1C) or "deactivated" (LF/HF > 10, green curve, Fig. 1C). We plotted LF/HF over time (Fig. 1D), showing that the example epochs highlighted in Fig. 1B were selected from periods of low LF/HF (black) and high LF/HF (green). Instead of utilizing LF/HF directly, we defined a transformed quantity $r = \log_{10}\frac{LF}{HF} - 1$, which is greater than 0 when LF/HF > 10. With this log-transform, ratio changes are additive rather than multiplicative.

Whisking activity was quantified from video recordings by extracting regions of interest around the whiskers and monitoring the frame-to-frame differences (Methods), which is collapsed into a single number quantifying whisker motion in arbitrary units (Fig. 1E). We found that r was negatively correlated with whisking activity (Fig. 1E; linear correlation between r and whisking activity, -0.43, p=10^{-22}; across N = 6 recordings, -0.35 ± 0.07), as expected from past findings (Fernandez et al., 2016) showing that high whisking activity is generally associated with low LF/HF.

We next examined the distribution of the cortical state metric ($P(r)$) during spontaneous activity throughout the experiment. We found that the distribution $P(r)$ (Fig. 1F; see Methods) was well modeled as a mixture of Gaussian distributions,

$$P(r) = \sum_{i=1}^{N} P(r|s^{\text{LFP}} = i)\, P(s^{\text{LFP}} = i),$$

where $P(r|s^{\text{LFP}} = i) \sim \mathcal{N}(\mu_i, \sigma_i^2)$ and $P(s^{\text{LFP}})$ is the probability of the system being in state $s^{\text{LFP}} \in \{1, \dots N\}$ (Fig. 1F). For the example in Fig. 1F, the model was implemented as a mixture of two distributions ($N = 2$), where the order was determined from the Bayesian Information Criterion

(BIC, Fig. 1G). Differences in BIC arise from the parameter cost term (which increases as more components are added) and the negative log-likelihood of the data (which becomes more negative as more components are added). Importantly, the two-component model outperformed the one-component model dramatically. Although this procedure is essentially curve-fitting, implicit is the concept that each of these components represents a particular state of the network, and the corresponding Gaussian is the conditional distribution $P(r|s^{\text{LFP}})$ for $s^{\text{LFP}} = 1$ (Fig. 1F, black) and $s^{\text{LFP}} = 2$ (Fig. 1F, green), while the marginal distribution $P(r)$ is the weighted mixture (Fig. 1F, black dots +/- error bars derived from counting statistics, with model overlaid in orange).

To examine the extent to which all individual recordings were consistent with the exemplar, we used the same mixture model approach on five additional recordings (Fig. 2). As an example of the distributions and models being fitted, Figure 2A and 2B show mixture models and empirical distributions of r for the best model (Fig. 2A) and the two-component model (Fig. 2B). For each recording, we plot in Figure 2C the parameters of each Gaussian component $(\mu \pm \sigma)$ fit to the data, along with the component mixing fractions (Fig. 2D), reflecting the amount of time spent in each state (i.e. $P(s^{\text{LFP}})$). As a guide to the eye, components with $\mu > 0$ are green and components with $\mu < 0$ are black or blue. In all recordings, there was a state with a mean r value that was less than 0 ("activated state", black and gray bars, Fig. 2D). The component weighted most heavily in the mixture model had mean (μ_i) less than 0 (Fig. 2D). While the recordings shared similarities, the precise values of the mean of each component varied across experiments. For example, in four recordings (1-1, 3-1, 3-2, and 4-1), the component parameters for a two-state model fell in a narrow range $(\mu_1 \in (-0.24, -0.14), \mu_2 \in (0.11, 0.15))$, while parameters for recording 2-1 were lower $(\mu_1 = -0.31, \mu_2 = 0.01)$.

Figure 2. LFP states mixture model parameters across recordings.
A: Mixture models fit using optimal number of components (based on BIC). Note that the vertical axis is log-scale, so each parabola (green, black, gray, blue) represents a component of the mixture distribution. Color in (A) corresponds to colors in (C) and (D). Recording 1-1 was shown in Fig. 1F.
B: Mixture models fit using two components.
C: Parameters of components in each mixture model as additional components are added, from 1 to 5 components.
D: Fraction of overall recording time allocated to each component.
E: Bayesian information criterion (BIC) curves for each recording. The minimum indicates where the information gained by high model fit quality is outweighed by the complexity of adding more parameters.

The best-fit model was determined from the BIC curves (Fig. 2E). In five of six recordings, the BIC analysis indicated more than one component was required to model the distribution of r. In two cases (Recording 2-1 and 3-1), the BIC analysis suggests a 3-component rather than 2-component fit, but the improvement was small, and there was no striking difference in the quality of fits shown in Fig. 2A-B. We therefore focus on the two-component LFP state models for the convenience of comparing across recordings. In the next several figures, we define state L ("low LF/HF") and state H ("high LF/HF") based on the heuristic threshold of LF/HF = 10 (equivalently, r = 0).

To summarize, an analysis of the statistics of a commonly used metric of cortical state shows that this metric is well-modeled by a mixture model with a small number of components. Further, we found that the distributions were overlapping, making a sharp distinction between states based on measuring r at a single time point challenging. Spiking activity also indicates cortical state, and numerous studies have shown how single neurons are correlated with certain aspects of the LFP (Cui *et al.*, 2016; Senzai, Fernandez-Ruiz and Buzsáki, 2019). However, how spontaneous neural *population* dynamics relate to the LFP has not been explored. Addressing this question directly is challenging due to the potentially high dimensionality of neural population activity. Thus, in the next sections, we model population spiking activity in a latent state framework and then analyze the relationship between spiking states and the LFP.

Latent state models of spiking activity

Simultaneously with the LFP recordings, we acquired spiking activity from an average of 26 neurons in each recording (range 9 to 32, N= 6 recordings, see Table 1; see Methods for spike

sorting metrics) distributed across layers 2 through 5 and categorized as FS (fast-spiking) and RS (regular-spiking) neurons as a function of their spike waveform. Due to variations in sampling across layers in each recording, we did not incorporate laminar location into the model of spiking dynamics. We observed that changes in spiking activity frequently occurred across a large fraction of the population simultaneously (Fig. 3A). To capture this structure, we fit a Poisson-emission hidden Markov model (Radons, Becket and Diilfer, 1994; Abeles *et al.*, 1995; Fontanini and Katz, 2008; Mazzucato, Fontanini and La Camera, 2015) to the population activity. We assume that neurons are independent, conditioned on the state, and that states evolve in time according to a Markov process. Specifically, the distribution of spike counts $\{y_n\}$ over a time bin of duration Δ ($\Delta = 40$ms) across the population ($n = 1 \ldots N_C$ neurons) while in spiking state s^{spk}, is given by

$$P(\{y_n\}|s^{\mathrm{spk}}) = \prod_{n=1}^{N_C} \frac{\left(\lambda_n\left(s^{\mathrm{spk}}\right)\Delta\right)^{k_j} e^{-\lambda_n(s^{\mathrm{spk}})\Delta}}{y_n!}$$

where each state $s^{\mathrm{spk}} \in \{1, 2, \ldots N_S\}$ is characterized by rate parameters $\lambda_n\left(s^{\mathrm{spk}}\right)$. The spiking state in the next time bin $t + 1$ depends only on the spiking state in the current time bin t, and is described by the transition probability matrix $A_{ij} = P(s^{\mathrm{spk}}(t + 1) = j|s^{\mathrm{spk}}(t) = i)$. Thus, a model with N_S states and N_C cells will have $N_S(N_S - 1) + N_S N_C$ parameters. Model parameters were estimated on half of the recording, and then a fixed model was used to infer state on the other half of the data (Methods). Importantly, although the individual neurons within the population are conditionally independent, their joint patterns of activity are determined by a common underlying state to capture the observation of simultaneous changes in activity across the diverse population.

Figure 3. Model of latent dynamics of spontaneous population spiking activity.
A: Spontaneous spiking activity recorded across 32 cells in awake mouse S1. Spikes from FS cells in blue; from RS cells in red.
B: States inferred by the HMM during the spontaneous segment in (A). Each row is a state and each column is a time bin. The shading (white to black) indicates the inferred probability of being in each state at each point in time.
C and D: Parameters of the HMM underlying inference in panel (B).
C: Firing rate parameters $(\lambda_n(s^{\text{spk}}))$ across the population of cells, including FS cells (blue) and RS cells (red). Bar length is the rate in spikes per 40-ms bin.
D: Matrix of state transition probabilities (A_{ij}). States in B, C, and D have been ordered based on average state-dependent firing rate across all cells.

To illustrate, we show a fitted model and inferred states in Figure 3, choosing the seven-state model ($N_S = 7$) for this recording (Recording 1-1 in Fig. 2) to illustrate the model fit, though we will return to the issue of choosing the number of states. Spontaneous spiking activity ($y_n(t)$) in a population of FS (blue) and RS (red) neurons is shown in the raster plot in Figure 3A. We compute the probability that the population is in a particular state at each point in time, conditioned on the entire observed sequence ($P(s^{\text{spk}}(t)|\{y_n(1:T)\})$, from $t = 1$ to $t = T$), shown in Fig. 3B. The model parameters, estimated on the training set consisting of half of the recording (Methods), are the rate parameters $\lambda_n(s^{\text{spk}})$ for each state s^{spk} and single unit n (Fig.

3C) and the transition matrix A_{ij} between states (Fig. 3D). The states are ordered by the average firing rate across cells, and the cells are grouped with the FS neurons (blue) and RS neurons (red) together. The same ordering is used throughout Figure 3 and 4. To get a qualitative sense of the structure captured by the model, note the "state 6" segments (Fig. 3B) corresponding to the population raster activity shown in Fig. 3A: these are marked by higher rates in the FS units (cells 26 to 32) and a few of the RS units (Fig. 3C). Some spike-states have overall lower firing rates than others (e.g., state 1 with 0.13 spikes/bin versus state 7 with 0.54 spikes/bin), but the distinction between states also derives from differences in the pattern of firing, as quantified by the Bhattacharyya coefficient, a generalized measure of differences between multinomial distributions, which ranges from 0 (no overlap) to 1 (full overlap). For instance, state 4 and state 5 have similar average firing rates (0.29 spikes/bin and 0.36 spikes/bin, respectively), but are highly distinguishable based on spiking pattern (Bhattacharyya coefficient of $2.8 \cdot 10^{-13}$; see Methods).

To assess how well the model fit the data, we generated a series of simulations of spontaneous S1 activity (Fig. 4). Using models ($A_{ij}, \lambda_n(s^{\mathrm{spk}})$, with 1 to 15 states) that were fit to the data as previously described, we simulated observations, matched to the data length of the original recording. Qualitatively, we see that a one-state model (Fig. 4A) does not show coordinated fluctuations in activity observed across multiple single units, as observed in Fig. 3A. A seven-state model (Fig. 3C, D) exhibits these transitions, visible in the simulated population activity shown in Fig. 4B. Note that the transitions in Fig. 4B are not expected to align with the transitions in the experimental data (Fig. 3A), as they were generated probabilistically from the model simulation. In the simulation, the sequence of states (simulation ground truth, Fig. 4C, top) was known, and were accurately inferred from the

Figure 4. Spiking state model validation metrics.
A: One-state model simulation for 32 cells for 10 s.
B: Seven-state model simulation using model fit (parameters shown in Figures 3C and 3D). A sequence of states was generated from the transition matrix (Figure 3D) and observations drawn according to the emission distributions (Poisson spiking, with rates shown in Figure 3C).
C: Top: the ground-truth state sequence used to generate spiking pattern (B). Bottom: inferred state probabilities, where dark purple indicates high probability.
D: Log-likelihood of test-set (reserved) data under models fit with 1 to 15 states (filled black circles). Log-likelihood of simulated data, using the fitted model and matching data length (open purple symbols).

simulated observations (inferred states, Fig. 4C, bottom). Thus, qualitatively, the multi-state Poisson-HMM captures the structure across single units and across time observed in the recorded data, and we now turn to quantitative measures of model fit.

To quantitatively assess model fits, we compared likelihood of decoded states in the data to that from ideal simulations and null tests generated by shuffling data. We first compared the likelihood of observations from simulated spontaneous S1 spiking (from the known model) to that of the recorded data (using a model fit on the reserved training set). Simulation length was matched to that of the recorded data (see Methods). We plot the log-likelihood (normalized by the

Figure 5. LFP states are coordinated with distinct states of spiking activity.
A: Example period of spontaneous activity in which LFP state transitioned from state H to state L.
B - C: Inferred spiking states (colors) with overlaid spike raster during segment 1 (LFP state H, panel B) and segment 2 (LFP state L, panel C). White indicates state probability was less than 80%.
D-E: Distribution of spiking states observed during LFP state H (panel D) and during LFP state L (panel E). Spiking states 1 and 2 are associated with LFP state H, while spiking states 5, 6, and 7 are associated with LFP state L.
F: The transition matrix (reproduced from Fig. 3D) was re-ordered using r_{LS}, the LFP state ratio (Methods), revealing block-diagonal structure.
G: Common transitions ($A_{ij} > 0.05$) between states (circles, color as in B and C) are indicated by directed arrows between nodes (states) in the graph. Self-connections (A_{ii}) not shown.

number of time bins times the number of neurons) for data (filled black circles) and simulation

(open purple circles) as a function of the number of states in the model (Fig. 4D). For this

recording, the log-likelihood of the data reaches that of the simulated data for models with five

states and more, indicating a fit that is as good as if the model were the "true" generating model for the data. Log-likelihood calculations for data, simulations, and shuffles for all recordings are provided in Figure 4-1. For all recordings, a multi-state model had higher likelihood than either the 1-state model or the shuffled model. However, it was not possible to identify a clear best-choice for the number of states based on cross-validated likelihood estimate (Fig. 4D). Thus, in subsequent figures, we continue to show examples from the seven-state model, as after this point the gain from additional states is minimal. All subsequent analyses were carried out for models fit with 2 to 10 states, with results shown in summary for different numbers of states.

Relationship of LFP and spiking states

With inferred spiking states calculated for each recording, we next asked how spiking states were related to LFP states (Figure 5). Within a single recording, there are periods of HF-dominated activity (low LF/HF, black, "state H") and LF-dominated activity (high LF/HF, green, "state L"), an example of which is shown in Figure 5A. The spiking states decoded during a segment of LFP state H ($r < 0$, Fig. 5B;) and LFP state L ($r > 0$, Fig. 5C) are shown with the population spike raster overlaid. Colors (gradient from black to green) represent distinct spiking states during the recording intervals for which spike-state was assigned with >80% probability. This example illustrates a tendency toward observing spiking states 1 and 2 during LFP state H (Fig. 5B) and a tendency toward observing spiking states 5, 6, and 7 during LFP state L (Fig. 5C). Across the full recording, we computed the relative frequency of spiking states observed during LFP state H or state L (Fig. 5D-E), and we defined, for each spiking state, the *LFP state ratio* (r_{LS}): the ratio of time that the LFP indicates state H ($r < 0$) to the time that LFP indicates state L ($r > 0$), conditioned on the identity of the inferred spiking state. Graphically, this quantity is the ratio of the heights of pairs of bars with the same color across Fig. 5D and Fig.

5E (Methods). We chose the ordering of states in Figure 5 based on r_{LS}. This ordering differs from that used in Figure 3, in which states were sorted by average firing rate, which produced a transition matrix with many off-diagonal elements (Fig. 5F, left). In contrast, sorting by r_{LS} revealed near-block diagonal structure in the transition matrix (Fig. 5F, right). We visualized the transition matrix as a graph (Fig. 5G), where the thickness of the arrows between states represents the probability of a transition. Graph nodes, representing spiking states, are colored as in panels D-E. Spiking states 1 and 2 [LFP state H] are close to each other and distant from spiking states 5, 6, and 7 [LFP state L]. Thus, while traversing the graph according to the transition matrix probabilities, the system will tend to stay in LFP-H spiking states or LFP-L spiking states for long periods of time.

Based on the tendency for spiking states to align with one LFP state or the other, we expected that using the state representation of spiking activity would effectively predict the LFP state (H or L). Thus, we built a binary classifier (Fig. 6) to predict the LFP state (H or L) based on either direct spiking metrics (single unit spike count, summed population spike count, or population vector) or on the inferred spiking states. Because LF/HF required a 1-second window to compute, for this analysis 40-ms time bins were sampled at 1 Hz across the experiment,

Figure 6. Decoding LFP state from spiking activity.
A. Decoding accuracy (chance level, 50%) of determining LFP state (H or L) from inferred spiking states (purple), population vector (solid black), and summed population rate (dashed) for an example recording (1-1).
B. Decoding accuracy, using single neuron activity (FS, left/blue; RS right/red), across all neurons.
C. Across 5 of 6 recordings, decoding accuracy is higher from inferred spiking states than from the population vector.

then divided randomly into training and test sets, balancing the number of observations associated with each LFP state so that chance-level performance was 50%.

For decoding, we used an SVM classifier with a linear kernel (see Methods), and we built a set of population activity decoders. In the simplest decoder, the predictor was the summed spike counts across the full population (single recording example, "summed population rate", Fig. 6A, dashed black line), which was a poor predictor of LFP state (accuracy, 54% ± 5%, range across 6 recordings 46% to 59%). To utilize the structure in the population firing pattern, we next constructed a linear classifier based on the population rate vector. The population vector classifier performed better than the summed-activity classifier (Fig. 6A, "population vector", solid black line), with performance in this example recording of 66% ± 3% (across cross-validation folds). Finally, we took as input the inferred spiking state and predicted LFP state as a function of the number of spiking states (Fig. 6A, "spiking states", solid purple line). This resulted in the highest decoding performance for models with at least 4 states. Decoding performance was comparable for models with 6 to 8 states, and then decreased, indicating little gain of constructing more complex models (Fig. 6A). To determine whether the accuracy of decoding LFP state was explained by a minority of high-performing cells, we used as predictors spike counts from single FS or RS neurons (Fig. 6B). While a few FS neurons were informative of LFP state (accuracy of best FS in each recording, 58 ± 7%, range 53% to 70%), most neurons decoded LFP state poorly (average decoding: 50% ± 2% (RS), N = 119; 52% ± 4% (FS), N = 39) (Fig. 6B). Across five of six recordings, the best-performing decoders used inferred spiking states (Fig. 6C). Decoding based on inferred spiking states substantially improved the ability to infer LFP state from spiking activity, achieving an average accuracy of 71% ± 10% (n = 6 recordings, range 61% to 85%). Figure 6C shows the decoding accuracy

using 7 spiking states. In comparison, decoding with the population vector was lower by 14 percentage points (n = 6 recordings, 57% ± 7%).

Spiking states are predictive of trial-by-trial variability in sensory-evoked responses

In the previous section, we found a consistent relationship between spiking states and the LFP. We now examine the extent to which latent spiking states are predictive of trial-by-trial variability in sensory-evoked responses. Between periods of spontaneous activity, we recorded sensory-evoked responses to simple punctate single-whisker stimuli as described previously (Sederberg *et al.*, 2019; Pala and Stanley, 2022), and we now analyze sensory-evoked responses in the latent states framework.

We first established that pre-stimulus state, classified under the spiking state model developed for spontaneous activity (Fig. 3), was predictive of single-trial evoked responses. We quantified the sensory-evoked responses, conditioned on the pre-stimulus state, and identified neurons with state-dependent sensory-evoked responses at a statistically significant level (Methods). Briefly, using one second (1 s) of pre-stimulus activity, we classified each trial based on the pre-stimulus spiking state that was the most probable state immediately prior to stimulus delivery. Trials were partitioned based on pre-stimulus spiking states, with population rate parameters plotted in Fig. 7A and sensory responses plotted in Fig. 7B. For comparison, the peri-stimulus time histogram (PSTH) across all trials is shown in gray (Fig. 7B). Qualitatively, in some states (state 1, Fig. 7A,B) the sensory-evoked response was substantially smaller than in other states (e.g., state 4, 5), both in terms of the absolute number of spikes in the post-stimulus window and the size of the evoked (i.e., subtracting the pre-stimulus spike count) response (Fig. 7B).

Figure 7. Spiking states are predictive of variability in sensory-evoked responses in a fraction of the population of single neurons.

A: Patterns of firing rates across neurons in pre-stimulus state for two recordings. Bars represent the model rate parameter (spikes per second) in single neurons in a 40-ms bin (as in Fig. 3). Color indicates r_{LS} as in Figure 3.

B: Responses to whisker stimulation (punctate deflection at time indicated by gray line) when trials are sorted by the decoded pre-stimulus state. The single neuron shown as an example in B is marked in panel A. PSTH bin width is 5 ms. Significant state-dependence is indicated by a purple star (facilitated) or yellow star (suppressed). This is defined relative to the overall average evoked response (gray bars) based on p< 0.01 in a 2-sided shuffle test (Methods).

C: Distribution of percentile values for state-averaged responses for all FS neurons for pre-stimulus-sorted (color) and trial-shuffled (gray) responses. Black dashed line indicates a uniform distribution, the theoretical value for the p-value distribution in the shuffle test which would be realized in the limit of infinite data ("shuffle (exact)").

D: Number of FS neurons with detected state-dependent responses using models with increasing numbers of states. Yellow/purple indicate suppressed/facilitated.

E: Same as C, for RS neurons.

F: Same as D, for RS neurons.

G: State transition graph (as in Figure 5), with purple/yellow indicating presence of facilitated / suppressed evoked responses in that state. Thicker outline represents more neurons with significant state dependence.

H: Cumulative count of the r_{LS} values of states in which state-dependent responses were identified in single neurons. Suppressed responses tended to occur in states more associated with LFP state H (low r_{LS}). Note that totals do not match the neuron counts in D and F because some individual neurons had facilitated (or suppressed) responses in more than one state.

We quantified statistical significance of the state dependence in sensory-evoked responses with a shuffle-based test which could accommodate the unequal number of trials assigned to each pre-stimulus state. Briefly, for each neuron and each state, we calculated the average evoked response across the set of trials assigned to the specified pre-stimulus state. We then generated a bootstrapped distribution, by repeating this computation on sets of randomly drawn trials, matching the number of trials per condition. We then computed the percentile of the state-dependent average evoked response within the boot-strapped distribution, and we show the distribution of these percentile values across all cells and all pre-stimulus states (Fig. 7C, E, "data", shown for FS neurons (7C) and RS neurons (7E) in a 7-state model). Based on this percentile value, state-conditioned responses were categorized as "suppressed" (<0.5; yellow star) or "facilitated" (>99.5; purple star), or not significantly different from the distribution of responses across all trials. We then repeated the analysis (for all cells and states) in which trial counts per state are matched but trial labels are shuffled to verify that the detected state-dependence was not simply a consequence of limited trial numbers (gray, "trial-shuffled", Fig. 7D, F). We found that more than half of neurons (for 7-state mode: 91 of 158 total, range 8 to 24 per recording, from N = 6 recordings; see Figure 7-1 for summary from the 7-state models) had at least one state-dependent (facilitated, 48 neurons; suppressed, 70 neurons; some cells are facilitated in one state and suppressed in another; Fig. 7C, E) sensory-evoked response, which exceeded the false-positive rate expected by our shuffle criterion (1%) or falsely detected in the shuffled controls (Fig. 7D, F; gray-shaded lines).

We next examined the organization of state-dependence across cells and across states. First, we found that state-modulated responses were more common in the FS population than the RS population. For instance, using the 7-state model, 58% of FS neurons were facilitated

(range: 37% to 80%, six recordings) and 70% of FS neurons were suppressed (range: 37% to 100%, six recordings), while 21% of RS neurons were facilitated (range: 9% to 31%) and 39% of RS neurons were suppressed (range: 23% to 54%). For this analysis, cells were only counted once; thus, if there were two states with suppressed responses (as in Fig. 7B), the neuron contributed once to the count in Fig. 7D. We note that, while we show the seven-state model for illustration, this choice was not critical for the detection of state-dependent responses. Models with four state or more show a relatively constant proportion of neurons with state-dependent responses (Fig. 7D, F).

Finally, we examined the extent to which specific pre-stimulus spiking states were associated with facilitated or suppressed responses. Based on a permutation test (Methods), we found that there was non-random structure of state-dependency in the evoked responses across the population, with a few specific states generating either facilitated or suppressed responses across multiple single units (see Figure 7-2 and Figure 7-3 for permutation test statistics). As in Fig. 5, we show state transition graphs, annotated to show the organization of facilitatory and suppressive pre-stimulus states. Rings around each state indicate the fraction of single units with either facilitated (Fig 7G, purple) or suppressed (Fig. 7G, yellow) responses, relative to the average evoked response across all trials, where thicker rings represent more neurons with significant state dependence. Suppressed states (yellow) were consistently found in the LFP-state-H associated spiking states (black), and a subset of LFP-state-L spiking states (lighter green) were associated with facilitated responses (purple). Across all recordings (Fig. 7H, using 7 spiking states), facilitated responses were most likely in spiking states with higher r_{LS} (LFP-state 2, median r_{LS}, 1.6), and suppressed responses were more likely during spiking states with lower r_{LS} (LFP-state 1, median r_{LS}, 0.6).

Spiking states are more informative of whisking activity than LF/HF

Finally, we examined the relationship between LFP, spiking states and whisking activity (Fig. 8). For a 30-second recording epoch, we show the LFP metric r (recall, $r = \log_{10} LF/HF - 1$, Fig. 8A), inferred spiking states (Fig. 8B), and whisking activity (Fig. 8C). 'Whisking activity' (Fig. 8C) was quantified as in Fig. 1. The example shown in Fig. 8 shows that higher whisking activity was observed when $r < 0$, as expected, but it also suggests that the spiking state representation may be more informative of whisking activity on a faster time scale. Spiking states 1 and 2 (dark blue shades) occur during whisking activity, with faster transitions that are, qualitatively, occurring on a similar timescale as the whisking activity fluctuations. To quantify this observation, we calculated the mutual information (MI) between whisking activity and the two different measures of state: r, based on LFP power ratio, and spiking states (Methods).

For the example recording in Fig. 8D, we found higher information rate using spiking states (0.25±0.02 bits/bin, whisking-spike states, bin of 40 ms) than direct thresholding of LF/HF (0.15 ± 0.02 bits/bin). In all recordings, using the spiking state representation (with at least 4 states, showing 4, 7, and 10 states in Fig. 8E, "whisk - spiking states") was more informative of whisking activity than thresholding LF/HF ("whisk – LFP"). We conclude that modeling the latent dynamics of spontaneous population spiking activity provides a useful metric of cortical state in awake mouse S1, containing information about changes in behavior with high temporal resolution.

Figure 8. Spiking states are more informative of whisking activity than LF/HF.

A: LFP state metric ($\log_{10}$ (LF/HF) − 1, defined in Fig. 1) during spontaneous activity. Gray shaded regions show where LF/HF < 10 and white for LF/HF > 10.

B: Spiking activity across 15 neurons and the inferred latent state during the same period of spontaneous activity. Colormap (differing from Fig. 5, 7) was chosen for higher contrast between states.

C: Whisking activity, quantified from motion energy (frame-to-frame changes, see Methods) in videography.

D: For this example, information between spiking state representation and whisking activity, for state models with 2 to 10 states. Information carried by comparing LF/HF to 10 is marked by the dotted black line.

E: Information about whisking activity calculated on all six recordings for 4-state, 7-state, and 10-state models are all higher than the information carried by LF/HF > 10. Error bars are estimated by resampling half of the data and repeating information calculation (50 repeats) and are typically smaller than the markers.

Discussion (1500 words)

A key aim of this work was to construct a bridge from classical notions of cortical state based on LFP to the activity of local ensembles of individual neurons. We have shown that latent states capture spatio-temporal structure in population spiking activity and are closely linked to slow-changing LFP states, predictive of single-trial sensory-evoked variability, and more informative of whisking activity than LFP. This framework is tailored to individual recordings, enabling future investigations of individual-specific differences in cortical state that may be informative of variations in perception and behavior.

Latent dynamics of population spiking activity in mouse S1

First, the latent-state representation of population activity inferred using the Poisson-HMM was highly informative of the "state" of the LFP. In fact, our analysis (Fig. 6) showed that the LFP states (based on LF/HF) were more accurately decoded from the latent state representation of spiking activity than they were from more direct measures, such as the summed firing rate across the population or even the population rate vector. It is not surprising that using the population without pooling across units provides more information: the population rate vector decoder outperformed the summed population activity decoder as well. However, the decoder using inferred spiking states also outperformed the population rate vector. The improvement of the inferred spiking states from the HMM over the population rate vector arises due to the temporal information embedded in the inferred spiking state, pointing to the importance of modeling correlations in time (i.e. dynamics) for analysis of spiking data. Thus, the Poisson-HMM framework provides a useful representation of population dynamics, capturing enough of the spatial (across cells) and temporal (across time) structure in cortical activity to enable more accurate decoding of the LFP state. While we have focused here on relating spiking states to

basic features of the LFP, we expect that the problem could be inverted. That is, a model based on an expanded representation of the LFP (e.g., using wavelets, LFP phase, or layer-specific features) could be used to infer latent spiking states.

Second, we showed that the inferred spiking states were predictive of single-trial sensory-evoked responses. In a previous study, we built a framework for predicting single-trial sensory-evoked LFP responses in awake mouse S1, in which features of the pre-stimulus LFP recording were identified based on their ability to predict response characteristics, rather than predetermined assumptions of informative features (Sederberg *et al.*, 2019). In that study, we found that our most predictive LFP features were equivalent to LF/HF, which had been used previously as a metric of state in S1 (Fernandez *et al.*, 2016; Poulet and Crochet, 2019). Thus, in the current study, we expected that pre-stimulus spiking states would also be predictive of single-trial sensory-evoked responses in single units, and that is what we found (Fig. 7). Indeed, the tendency of so-called "suppressed response" pre-stimulus states to align with low LF/HF is highly consistent with the predictive features examined in our previous study. In addition to LF/HF, our previous work used the phase (or activation/deactivation of LFP) to predict sensory-evoked responses. Here we did not examine the relationship between LFP phase and spiking states, but we would expect LFP phase to align closely with spiking states associated with high LF/HF ("green states" in Fig. 7).

Our results also uncovered two functional relationships captured by latent dynamical states. First, we showed that predictive pre-stimulus states tended to result in either facilitatory or suppressive influence on evoked responses, similar to previous studies in V1 showing that a shared gain and offset can account for population-wide patterns of variability (Lin *et al.*, 2015). However, in our study, we also find that sensory responses in a substantial fraction of the

population were *not* modulated by state. Interestingly, we found that FS neurons, which are more likely to be inhibitory cells (Barthó *et al.*, 2004), were more likely than RS neurons to exhibit state-dependent sensory-evoked response patterns (Fig. 7). Second, we found that spiking states were informative of whisking activity, consistent with other recent studies predicting spontaneous neural activity from monitored mouse behaviors (Musall *et al.*, 2019; Stringer *et al.*, 2019; Salkoff *et al.*, 2020). Functional relationships like these (i.e. modulation by pre-stimulus state, relationship to whisking) are informative of anatomical cell type (Gentet *et al.*, 2010, 2012; Muñoz *et al.*, 2017; Yu *et al.*, 2019) and thus expand the set of measures available for cell type classification in extracellular recordings.. How specific cell types across cortical layers shape latent dynamics could be investigated using new datasets that provide more specific cell-type identity information (de Vries *et al.*, 2020).

Relationship to other studies using the Poisson-HMM to analyze spiking activity.

In the Poisson-HMM used here to model spiking states, all cells were assumed to be linked to a single hidden state at any moment of time, and this state is assumed to have Markovian dynamics, with dependence only on the current state. These assumptions limit the flexibility of the model, but even so, inferred spiking states from the model were successful in predicting LFP state, sensory-evoked responses, and whisking activity. Several previous studies have used the Poisson-HMM framework to describe population activity patterns and dynamics (Jones *et al.*, 2007; Mazzucato, Fontanini and La Camera, 2015; Engel *et al.*, 2016; Prentice *et al.*, 2016). For instance, the Poisson-HMM captured the structure of fast sequences of population activity patterns in rat gustatory cortex, demonstrating step-like rather than ramping behavior during a sensory-decision-making task (Jones *et al.*, 2007). We used a coarser time bin, lasting 40 ms rather than 1 – 2 ms in the Jones study. We chose this bin size because our analysis concerned spontaneous activity over relatively long (minutes) segments of data. Thus,

while we use the same kind of model, it is simply for the expedience of generating a relatively simple representation of spatial and temporal dynamics. The same model was also used to capture spontaneous activity states in the five seconds preceding a sensory-evoked response (Mazzucato, Fontanini and La Camera, 2015), and there it was argued that the identified states reflect underlying attractor dynamics. Again, due to the prevalence of FS units in the spiking patterns that determined spiking states, a different interpretation of states may be more appropriate in our study. For instance, a recent modeling study showed that PV interneurons are poised to switch the cortical network in and out of an inhibition-stabilized network (ISN) dynamical state (Bryson *et al.*, 2021). Such state switches would be reflected in differences in sensory-evoked responses to well-controlled stimuli, as we observed, and it would be interesting to investigate the extent to which ISN dynamics predict the specific patterns of evoked sensory responses we observed.

Limitations of Poisson-HMM for spiking dynamics

While the Poisson-HMM has the advantage of a small number of parameters and simple state dynamics, the assumption that all neurons are in the same state likely breaks down in populations spanning larger regions of cortex. An alternative approach, more commonly used in non-sensory brain areas (motor, prefrontal, hippocampus) is to embed neuronal activity patterns in a low-dimensional manifold upon which population activity smoothly evolves. For example, latent structure in population dynamics inferred utilizing a dynamical system improved decoding of movement intention (Pandarinath *et al.*, 2015), and other studies have identified low-dimensional structure in prefrontal cortex (Mante *et al.*, 2013) and hippocampus (Nieh *et al.*, 2021). It is not clear to what extent the manifold framework generalizes to sensory pathways. For instance, responses in visual cortex to sensory inputs appear to live in a distinct subspace from ongoing, spontaneous activity (Stringer *et al.*, 2019), but this is a statement

about spatial correlations, not about dynamics. This issue could be clarified by developing more rigorous "null" computational models, where latent dynamics and population structure can be carefully controlled (Sederberg and Nemenman, 2020; Morrell, Sederberg and Nemenman, 2021).

Future opportunities for data-driven cortical state investigation

Understanding of the dynamics of local brain state has advanced tremendously through careful dissection of modulatory circuit mechanisms in primary sensory areas. Experiments utilizing patch-clamp electrophysiology in awake mice have revealed state-dependent changes in the activity of single neurons of specific cell classes (Gentet *et al.*, 2010; Pala and Petersen, 2018), shared patterns of variability across small numbers of cells (Arroyo, Bennett and Hestrin, 2018), and control of excitability through specific cell type activation (Fu *et al.*, 2014). In the large-scale multi-single-unit recordings enabled by modern high-density multielectrode recordings, such fine-scale information is rarely available. On the other hand, highly parallel recording methods reveal functional relationships across the population that are not observable from recordings of single neurons, and likely carry rich information on the functional roles of specific anatomical cell types. The latent dynamics uncovered by these approaches may also provide functionally meaningful and pragmatic targets for control of circuit and network activity as tools for applying control theory to neural circuits continue to emerge (Bolus *et al.*, 2021). Better understanding, through quantitative frameworks such as that presented here, of these features of cortical dynamics will ultimately provide an invaluable tool to map individual-specific differences in the neural signatures of behavior, perception, and cognition.

References

Abeles, M. *et al.* (1995) 'Cortical activity flips among quasi-stationary states.', *Proceedings of the National Academy of Sciences*, 92(19), pp. 8616–8620. doi: 10.1073/pnas.92.19.8616.

Arroyo, S., Bennett, C. and Hestrin, S. (2018) 'Correlation of Synaptic Inputs in the Visual Cortex of Awake, Behaving Mice', *Neuron*. Elsevier Inc., 99(6), pp. 1289-1301.e2. doi: 10.1016/j.neuron.2018.08.008.

Barthó, P. *et al.* (2004) 'Characterization of Neocortical Principal Cells and Interneurons by Network Interactions and Extracellular Features', *Journal of Neurophysiology*, 92(1), pp. 600–608. doi: 10.1152/jn.01170.2003.

Beaman, C. B., Eagleman, S. L. and Dragoi, V. (2017) 'Sensory coding accuracy and perceptual performance are improved during the desynchronized cortical state', *Nature Communications*. Springer US, 8(1), pp. 1–14. doi: 10.1038/s41467-017-01030-4.

Bennett, C., Arroyo, S. and Hestrin, S. (2013) 'Subthreshold Mechanisms Underlying State-Dependent Modulation of Visual Responses', *Neuron*. Elsevier Inc., 80(2), pp. 350–357. doi: 10.1016/j.neuron.2013.08.007.

Bhattacharyya, A. (1946) 'On a Measure of Divergence between Two Multinomial Populations', *Sankhya: The Indian Journal of Statistics*, 7(4), pp. 401–406. Available at: www.jstor.org/stable/25047882.

Bolus, M. F. *et al.* (2021) 'State-space optimal feedback control of optogenetically driven neural activity', *Journal of Neural Engineering*, 18(3). doi: 10.1088/1741-2552/abb89c.

Bryson, A. *et al.* (2021) 'State transitions through inhibitory interneurons in a cortical network model', *PLoS Computational Biology*, 17(10 October), pp. 1–29. doi: 10.1371/journal.pcbi.1009521.

Cui, Y. *et al.* (2016) 'Inferring Cortical Variability from Local Field Potentials', *The Journal of Neuroscience*, 36(14), pp. 4121–4135. doi: 10.1523/JNEUROSCI.2502-15.2016.

Eggermann, E. *et al.* (2014) 'Cholinergic Signals in Mouse Barrel Cortex during Active Whisker Sensing', *Cell Reports*. The Authors, 9(5), pp. 1654–1660. doi: 10.1016/j.celrep.2014.11.005.

Engel, T. A. *et al.* (2016) 'Selective modulation of cortical state during spatial attention', *Science*, 354(6316), pp. 1140–1144.

Fernandez, L. M. J. *et al.* (2016) 'Highly Dynamic Spatiotemporal Organization of Low-Frequency Activities During Behavioral States in the Mouse Cerebral Cortex', *Cerebral Cortex*, (January 2018), pp. 5444–5462. doi: 10.1093/cercor/bhw311.

Fontanini, A. and Katz, D. B. (2008) 'Behavioral States, Network States, and Sensory Response Variability', *Journal of Neurophysiology*, 100(3), pp. 1160–1168. doi: 10.1152/jn.90592.2008.

Fu, Y. *et al.* (2014) 'A Cortical Circuit for Gain Control by Behavioral State', pp. 1139–1152. doi: 10.1016/j.cell.2014.01.050.

Gentet, L. J. *et al.* (2010) 'Membrane Potential Dynamics of GABAergic Neurons in the Barrel Cortex of Behaving Mice', *Neuron*. Elsevier Ltd, 65(3), pp. 422–435. doi: 10.1016/j.neuron.2010.01.006.

Gentet, L. J. *et al.* (2012) 'Unique functional properties of somatostatin-expressing GABAergic neurons in mouse barrel cortex', *Nature Neuroscience*. Nature Publishing Group, 15(4), pp. 607–612. doi: 10.1038/nn.3051.

Gratiy, S. L. *et al.* (2011) 'On the estimation of population-specific synaptic currents from laminar multielectrode recordings', 5(December), pp. 1–19. doi: 10.3389/fninf.2011.00032.

Hagen, E. *et al.* (2016) 'Hybrid Scheme for Modeling Local Field Potentials from Point-Neuron Networks', (October), pp. 4461–4496. doi: 10.1093/cercor/bhw237.

Haslinger, R. *et al.* (2006) 'Analysis of LFP phase predicts sensory response of barrel cortex.', *Journal of Neurophysiology*, 96(3), pp. 1658–1663. doi: 10.1152/jn.01288.2005.

He, B. J. and Zempel, J. M. (2013) 'Average Is Optimal: An Inverted-U Relationship between Trial-to-Trial Brain Activity and Behavioral Performance', *PLoS Computational Biology*, 9(11). doi: 10.1371/journal.pcbi.1003348.

Jones, L. M. *et al.* (2007) 'Natural stimuli evoke dynamic sequences of states in sensory cortical ensembles', *Proceedings of the National Academy of Sciences of the United States of America*, 104(47), pp. 18772–18777. doi: 10.1073/pnas.0705546104.

Köster, U. *et al.* (2014) 'Modeling Higher-Order Correlations within Cortical Microcolumns', *PLoS Computational Biology*, 10(7), pp. 1–12. doi: 10.1371/journal.pcbi.1003684.

Lin, I. C. *et al.* (2015) 'The Nature of Shared Cortical Variability', *Neuron*. The Authors, 87(3), pp. 645–657. doi: 10.1016/j.neuron.2015.06.035.

Ludwig, K. A. *et al.* (2006) 'Chronic neural recordings using silicon microelectrode arrays electrochemically deposited with a poly(3,4-ethylenedioxythiophene) (PEDOT) film', *Journal of Neural Engineering*, 3(1), pp. 59–70. doi: 10.1088/1741-2560/3/1/007.

Mante, V. *et al.* (2013) 'Context-dependent computation by recurrent dynamics in prefrontal cortex', *Nature*. Nature Publishing Group, 503(7474), pp. 78–84. doi: 10.1038/nature12742.

Mazzucato, L., Fontanini, A. and La Camera, G. (2015) 'Dynamics of Multistable States during

Ongoing and Evoked Cortical Activity', *Journal of Neuroscience*, 35(21), pp. 8214–8231. doi: 10.1523/JNEUROSCI.4819-14.2015.

McCormick, D. A., Nestvogel, D. B. and He, B. J. (2020) 'Neuromodulation of Brain State and Behavior', *Annual Review of Neuroscience*, pp. 391–415. doi: 10.1146/annurev-neuro-100219-105424.

Monto, S. *et al.* (2008) 'Very Slow EEG Fluctuations Predict the Dynamics of Stimulus Detection and Oscillation Amplitudes in Humans', *Journal of Neuroscience*, 28(33), pp. 8268–8272. doi: 10.1523/JNEUROSCI.1910-08.2008.

Morrell, M. C., Sederberg, A. J. and Nemenman, I. (2021) 'Latent Dynamical Variables Produce Signatures of Spatiotemporal Criticality in Large Biological Systems', *Physical Review Letters*. American Physical Society, 126(11), p. 118302. doi: 10.1103/PhysRevLett.126.118302.

Muñoz, W. *et al.* (2017) 'Layer-specific modulation of neocortical dendritic inhibition during active wakefulness', *Science*, 959(March), pp. 954–959.

Musall, S. *et al.* (2019) 'Single-trial neural dynamics are dominated by richly varied movements', *Nature Neuroscience*. Springer US, 22(10), pp. 1677–1686. doi: 10.1038/s41593-019-0502-4.

Nieh, E. H. *et al.* (2021) 'Geometry of abstract learned knowledge in the hippocampus', *Nature*. Springer US, 595(7865), pp. 80–84. doi: 10.1038/s41586-021-03652-7.

Okun, M. *et al.* (2015) 'Diverse coupling of neurons to populations in sensory cortex', *Nature*, 521(7553), pp. 511–515. doi: 10.1038/nature14273.

Pachitariu, M. *et al.* (2016) 'Fast and accurate spike sorting of high-channel count probes with KiloSort', *Advances in Neural Information Processing Systems*, (Nips), pp. 4455–4463.

Pala, A. and Petersen, C. C. (2018) 'State-dependent cell-type-specific membrane potential dynamics and unitary synaptic inputs in awake mice', *eLife*, 7, pp. 1–13. doi: 10.7554/eLife.35869.

Pala, A. and Stanley, G. B. (2022) *Ipsilateral stimulus encoding in primary and secondary somatosensory cortex of awake mice*, *bioRxiv*. Available at: https://www.biorxiv.org/content/10.1101/2021.07.16.452704v1%0Ahttps://www.biorxiv.org/content/10.1101/2021.07.16.452704v1.abstract.

Pandarinath, C. *et al.* (2015) 'Neural population dynamics in human motor cortex during movements in people with ALS.', *eLife*, 4, p. e07436. doi: 10.7554/eLife.07436.

Petersen, C. C. H. *et al.* (2003) 'Interaction of sensory responses with spontaneous

depolarization in layer 2/3 barrel cortex', *Proceedings of the National Academy of Sciences of the United States of America*, 100(23), pp. 13638–13643. doi: 10.1073/pnas.2235811100.

Pinto, L. *et al.* (2013) 'Fast modulation of visual perception by basal forebrain cholinergic neurons', *Nature Neuroscience*. Nature Publishing Group, 16(12), pp. 1857–1863. doi: 10.1038/nn.3552.

Poulet, J. F. A. and Crochet, S. (2019) 'The cortical states of wakefulness', *Frontiers in Systems Neuroscience*, 12(January), pp. 1–18. doi: 10.3389/fnsys.2018.00064.

Poulet, J. F. A. and Petersen, C. C. H. (2008) 'Internal brain state regulates membrane potential synchrony in barrel cortex of behaving mice', *Nature*, 454(7206), pp. 881–885. doi: 10.1038/nature07150.

Prentice, J. S. *et al.* (2016) 'Error-Robust Modes of the Retinal Population Code', pp. 1–34. doi: 10.5061/dryad.1f1rc.Funding.

Radons, G., Becket, D. and Diilfer, B. (1994) 'Analysis , classification , and coding of multielectrode spike trains with hidden Markov models', 373, pp. 359–373.

Rossant, C. and Harris, K. D. (2013) 'Hardware-accelerated interactive data visualization for neuroscience in Python', *Frontiers in Neuroinformatics*, 7(DEC), pp. 1–9. doi: 10.3389/fninf.2013.00036.

Salkoff, D. B. *et al.* (2020) 'Movement and Performance Explain Widespread Cortical Activity in a Visual Detection Task', *Cerebral Cortex*, 30(1), pp. 421–437. doi: 10.1093/cercor/bhz206.

Sederberg, A. J. *et al.* (2019) 'State-aware detection of sensory stimuli in the cortex of the awake mouse', *PLOS Computational Biology*. Public Library of Science, 15(5), p. e1006716. doi: 10.1371/journal.pcbi.1006716.

Sederberg, A. and Nemenman, I. (2020) 'Randomly connected networks generate emergent selectivity and predict decoding properties of large populations of neurons', *PLOS Computational Biology*. Public Library of Science, 16(5), p. e1007875. Available at: https://doi.org/10.1371/journal.pcbi.1007875.

Senzai, Y., Fernandez-Ruiz, A. and Buzsáki, G. (2019) 'Layer-Specific Physiological Features and Interlaminar Interactions in the Primary Visual Cortex of the Mouse', *Neuron*, 101(3), pp. 500-513.e5. doi: 10.1016/j.neuron.2018.12.009.

Speed, A. *et al.* (2019) 'Cortical State Fluctuations across Layers of V1 during Visual Spatial Perception', *Cell Reports*. ElsevierCompany., 26(11), pp. 2868-2874.e3. doi: 10.1016/j.celrep.2019.02.045.

Stringer, C. *et al.* (2019) 'Spontaneous behaviors drive multidimensional, brainwide activity',

Science, 364(6437), p. eaav7893. doi: 10.1126/science.aav7893.

Vinck, M. *et al.* (2015) 'Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding Article Arousal and Locomotion Make Distinct Contributions to Cortical Activity Patterns and Visual Encoding', *Neuron*. Elsevier Inc., 86(3), pp. 740–754. doi: 10.1016/j.neuron.2015.03.028.

de Vries, S. E. J. *et al.* (2020) 'A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex', *Nature Neuroscience*, 23(1), pp. 138–151. doi: 10.1038/s41593-019-0550-9.

Wang, Q., Webber, R. M. and Stanley, G. B. (2010) 'Thalamic synchrony and the adaptive gating of information flow to cortex', *Nature Neuroscience*. Nature Publishing Group, 13(12), pp. 1534–1541. doi: 10.1038/nn.2670.

Yu, J. *et al.* (2019) 'Recruitment of GABAergic Interneurons in the Barrel Cortex during Active Tactile Behavior', *Neuron*. Elsevier Inc., 104(2), pp. 412-427.e4. doi: 10.1016/j.neuron.2019.07.027.

Figure Legends

Figure 1. LFP states during wakefulness.

A: Local field potential (LFP) and single unit (SU) activity were recorded from laminar array electrodes in primary somatosensory (S1) cortex of awake head-fixed mouse.

B: LFP in layer 4 over a 20-s period. Time axes in B, D, and I are the same. Black bar: example of *activated* state ("H"). Green bar: example of the *deactivated* state ("L").

C: Fourier amplitude of the LFP. LF is the average amplitude between 1 and 10 Hz and HF is the average amplitude between 30 and 90 Hz. Fourier amplitude spectrum averaged over activated (black) and inactivated (green) periods, where states were classified based on LF/HF ratio.

D: LF/HF evaluated in sliding 1-s window over 20 seconds, with black/green shading corresponding to bars in (B).

E: Density plot (color represents log density) of the cortical state metric ($r = log \, (LF/HF) - 1$) and whisker motion, quantified from video recording (see Methods). Linear correlation was -0.43 (p = 1e-22).

F: Distribution of r and fit to a mixture of Gaussians model with two components.

G: Bayesian Information Criterion (BIC) for mixture of Gaussians model for 1 to 8 states is minimized at n = 2 states.

Figure 2. Mixture model parameters across recordings.

A: Mixture models fit using optimal number of components (based on BIC). Note that the vertical axis is log-scale, so each parabola (green, black, gray, blue) represents a component of the mixture distribution. Color in (A) corresponds to colors in (C) and (D). Recording 1-1 was shown in Fig. 1F.

B: Mixture models fit using two components.

C: Parameters of components in each mixture model as additional components are added, from 1 to 5 components.

D: Fraction of overall recording time allocated to each component.

E: Bayesian information criterion (BIC) curves for each recording. The minimum indicates where the information gained by high model fit quality is outweighed by the complexity of adding more parameters.

Figure 3. Model of latent dynamics of spontaneous spiking activity.
A: Spontaneous spiking activity recorded across 32 cells in awake mouse S1. Spikes from FS cells in blue; from RS cells in red.

B: States inferred by the HMM during the spontaneous segment in (A). Each row is a state and each column is a time bin. The shading (white to black) indicates the inferred probability of being in each state at each point in time.

C and D: Parameters of the HMM underlying inference in panel (B).

C: Firing rate parameters $(\lambda_n(s^{\text{spk}}))$ across the population of cells, including FS cells (blue) and RS cells (red). Bar length is the rate in spikes per 40-ms bin.

D: Matrix of state transition probabilities (A_{ij}). States in B, C, and D have been ordered based on average state-dependent firing rate across all cells.

Figure 4. Spiking state model validation metrics.
A: "One-state model" simulation for 32 cells over 10 seconds.

B: Simulated data using model fit (parameters show in C and D). A sequence of states was generated from the transition matrix (D) and observations drawn according to the emission distributions (Poisson spiking, with rates shown in C).

C: Top: the ground-truth state sequence used to generate spiking pattern (B). Bottom: inferred state probabilities, where dark purple indicates high probability.

D: Log-likelihood of test-set (reserved) data under models fit with 1 to 15 states (filled black circles). Log-likelihood of simulated data, using the fitted model and matching data length (open purple symbols). See Figure 4-1 for log-likelihood values for all recordings.

Figure 5. LFP states are coordinated with distinct states of spiking activity.

A: Example period of spontaneous activity in which LFP state transitioned from state H to state L.

B - C: Inferred spiking states (colors) with overlaid spike raster during segment 1 (LFP state H, panel B) and segment 2 (LFP state L, panel C). White indicates state probability was less than 80%.

D-E: Distribution of spiking states observed during LFP state H (panel D) and during LFP state L (panel E). Spiking states 1 and 2 are associated with LFP state H, while spiking states 5, 6, and 7 are associated with LFP state L.

F: The transition matrix (reproduced from Fig. 3D) was re-ordered using r_{LS}, the LFP state ratio (Methods), revealing block-diagonal structure.

G: Common transitions ($A_{ij} > 0.05$) between states (circles, color as in B and C) are indicated by directed arrows between nodes (states) in the graph. Self-connections (A_{ii}) not shown.

Figure 6. Decoding LFP state from spiking activity.

A. Decoding accuracy (chance level, 50%) of determining LFP state (H or L) from inferred spiking states (purple), population vector (solid black), and summed population rate (dashed) for an example recording (1-1).

B. Decoding accuracy, using single neuron activity (FS, left/blue; RS right/red), across all neurons.

C. Across 5 of 6 recordings, decoding accuracy is higher from inferred spiking states than from the population vector.

Figure 7. Spiking states are predictive of variability in sensory-evoked responses in a fraction of the population of single neurons.

A: Patterns of firing rates across neurons in pre-stimulus state for two recordings. Bars represent the model rate parameter (spikes per second) in single neurons in a 40-ms bin (as in Fig. 3). Color indicates r_{LS} as in Figure 3.

B: Responses to whisker stimulation (punctate deflection at time indicated by gray line) when trials are sorted by the decoded pre-stimulus state. The single neuron shown as an example in B is marked in panel A. PSTH bin width is 5 ms. Significant state-dependence is indicated by a purple star (facilitated) or yellow star (suppressed). This is defined relative to the overall average evoked response (gray bars) based on p< 0.01 in a 2-sided shuffle test (Methods).

C: Distribution of percentile values for state-averaged responses for all FS neurons for pre-stimulus-sorted (color) and trial-shuffled (gray) responses. Black dashed line indicates a uniform distribution, the theoretical value for the p-value distribution in the shuffle test which would be realized in the limit of infinite data ("shuffle (exact)").

D: Number of FS neurons with detected state-dependent responses using models with increasing numbers of states. Yellow/purple indicate suppressed/facilitated.

E: Same as C, for RS neurons.

F: Same as D, for RS neurons.

G: State transition graph (as in Figure 5), with purple/yellow indicating presence of facilitated / suppressed evoked responses in that state. Thicker outline represents more neurons with significant state dependence.

H: Cumulative count of the r_{LS} values of states in which state-dependent responses were identified in single neurons. Suppressed responses tended to occur in states more associated

with LFP state H (low r_{LS}). Note that totals do not match the neuron counts in D and F because some individual neurons had facilitated (or suppressed) responses in more than one state.

Figure 8. Spiking states are more informative of whisking activity than LF/HF.

A: LFP state metric ($\log_{10}$ (LF/HF) − 1, defined in Fig. 1) during spontaneous activity. Gray shaded regions show where LF/HF < 10 and white for LF/HF > 10.

B: Spiking activity across 15 neurons and the inferred latent state during the same period of spontaneous activity. Colormap (differing from Fig. 5, 7) was chosen for higher contrast between states.

C: Whisking activity, quantified from motion energy (frame-to-frame changes, see Methods) in videography.

D: For this example, information between spiking state representation and whisking activity, for state models with 2 to 10 states. Information carried by comparing LF/HF to 10 is marked by the dotted black line.

E: Information about whisking activity calculated on all six recordings for 4-state, 7-state, and 10-state models are all higher than the information carried by LF/HF > 10. Error bars are estimated by resampling half of the data and repeating information calculation (50 repeats) and are typically smaller than the markers.

Extended Data and Extended Data Legends

Recording	Data (1-state model)	Data (7-state model)	Simulation (7-state)	Shuffled (7-state)
1-1	-428549	-399075	-401893	-437928
1-2	-295463	-279170	-269299	-287421
2-1	-426533	-400726	-401244	-431003
3-1	-255510	-220216	-217860	-252243
3-2	-508583	-469814	-448926	-500754
4-1	-198327	-185925	-184549	-200262

Figure 4-1. Log likelihood calculations under Poisson-HMM. Comparison of data (or simulated data) log-likelihoods under the Poisson-HMM, for the recorded data (1-state), recorded data (7-state), simulated data (7-state), shuffled simulated data (7-state). Simulation length is matched to recorded data length, so comparisons can be made within each row, but not directly across rows.

Recording (# SU)	7-state tests # facilitated	7-state tests # suppressed	F(+) response: pattern p-val	S(-) response: pattern p-val
1-1 (32)	12	14	2.3e-4	2.5e-5
1-2 (32)	9	12	2.6e-2	1.8e-2
2-1 (28)	5	8	1.5e-1	1.8e-4
3-1 (11)	7	12	5.4e-2	6.7e-2
3-2 (36)	19	34	1.8e-4	1.5e-3
4-1 (19)	9	15	2.7e-3	8.0e-3

Figure 7-1. Number of state-dependent evoked responses detected in each recording, under the 7-state model and results of test for pattern of state-dependent responses across single units within a recording. Reported p-value for response pattern detection is the fraction of

random evoked-response patterns that have as many facilitated responses (suppressed responses) grouped in a single state as found in the empirical results.

Recording	2 states	3 states	4 states	5 states	6 states	7 states	8 states
1-1	7.81e-03	2.59e-03	3.51e-03	1.66e-04	1.56e-04	2.50e-05	1.69e-04
1-2	6.25e-02	3.50e-01	2.35e-01	1.44e-02	9.58e-02	1.79e-02	3.13e-01
2-1	6.25e-02	1.78e-02	6.10e-05	7.73e-04	2.36e-04	1.83e-04	8.65e-04
3-1	1.09e-01	1.23e-02	2.54e-02	3.12e-02	2.47e-02	6.69e-02	3.18e-02
3-2	1.53e-05	2.13e-04	1.79e-02	2.24e-02	9.65e-03	1.48e-03	9.26e-03
4-1	6.25e-02	9.65e-04	2.30e-03	3.07e-03	2.91e-03	7.97e-03	4.59e-02

Figure 7-2. Permutation test for organization of suppressed sensory-evoked responses across single units for N = 2 to 8 states. Probability of observing clusters of SUPPRESSED responses by recording and number of states in the model. Gray indicates p > 0.05.

Recording	2 states	3 states	4 states	5 states	6 states	7 states	8 states
1-1	7.81e-03	2.59e-03	7.25e-04	4.17e-03	1.83e-03	2.26e-04	1.58e-04
1-2	3.12e-01	8.79e-02	3.91e-03	3.20e-04	4.37e-02	2.59e-02	4.63e-02
2-1	6.25e-02	2.59e-01	3.67e-01	2.72e-02	7.41e-02	1.52e-01	7.89e-02
3-1	3.12e-02	4.12e-03	7.81e-03	2.58e-01	7.41e-02	5.43e-02	1.63e-02
3-2	4.88e-04	2.85e-04	1.44e-03	9.77e-03	3.92e-07	1.77e-04	4.96e-06
4-1	6.25e-02	5.08e-05	1.53e-05	5.12e-06	1.59e-06	2.68e-03	5.34e-04

Figure 7-3. Permutation test for organization of facilitated sensory-evoked responses across single units for N = 2 to 8 states. Probability of observing clusters of FACILITATED responses by recording and number of states in the model. Gray indicates p > 0.05.