Neural manifolds in V1 change with top-down signals from V4 targeting the foveal region

Aitor Morales-Gregorio[1,2*], Anno C. Kurth[1,3], Junji Ito[1], Alexander Kleinjohann[1,4], Frédéric V. Barthélemy[1,5], Thomas Brochier[5], Sonja Grün[1,4], Sacha J. van Albada[1,2]

[1]Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA-Institut Brain Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany. [2]Institute of Zoology, University of Cologne, Cologne, Germany. [3]RWTH Aachen University, Aachen, Germany. [4]Theoretical Systems Neurobiology, RWTH Aachen University, Aachen, Germany. [5]Institut de Neurosciences de la Timone (INT), CNRS & Aix-Marseille Université, Marseille, France.

[*]Corresponding author: a.morales-gregorio@fz-juelich.com

June 14, 2023

High-dimensional brain activity is often organised into lower-dimensional neural manifolds. However, the neural manifolds of the visual cortex remain understudied. Here, we study large-scale multielectrode electrophysiological recordings of macaque (*Macaca mulatta*) areas V1, V4 and DP with a high spatio-temporal resolution. We find, for the first time, that the population activity of V1 contains two separate neural manifolds, which correlate strongly with eye closure (eyes open/closed) and have distinct dimensionalities. Moreover, we find strong top-down signals from V4 to V1, particularly to the foveal region of V1, which are significantly stronger during the eyes-open periods, a previously unknown effect. Finally, *in silico* simulations of a balanced spiking neuron network qualitatively reproduce the experimental findings. Taken together, our analyses and simulations suggest that top-down signals modulate the population activity of V1, causing two distinct neural manifolds. We postulate that the top-down modulation during the eyes-open periods prepares V1 for fast and efficient visual responses, resulting in a type of visual stand-by mode.

Keywords: Neural manifold, dimensionality, electrophysiology, macaque, visual cortex, top-down signals, resting state

Introduction

The brain can be described as a high-dimensional dynamical system capable of representing and processing a plethora of low-dimensional variables.

The time-resolved activity of a population of neurons can be considered as a trajectory in a high-dimensional space, where each neuron represents one dimension; i.e., the state space of the neural system. Typically, the system does not attain all possible states in the state space, but rather remains confined to small subsets. These subsets of the state space are referred to as a neural manifolds[1–5]. Neural manifolds have been shown to encode aspects such as decision-making in the prefrontal cortex of macaque[6], hand movement trajectories in the motor cortex of macaque[2,3,7], odour in the piriform cortex of mice[8], head direction in the anterodorsal thalamic nucleus of mice[5], and spatial position in the hippocampus of mice[9]. The study of neural manifolds in the visual cortex has been conducted in mice[10,11] and macaque[12,13]. However, to the best of our knowledge, the state dependence of neural manifolds in the primary visual cortex (V1) of macaque has not yet been investigated.

Neural manifolds often have an intricate structure, which can be studied using methods borrowed from computational topology[5,12,14]. In addition to the topology, the number of uncorrelated covariates required to capture the variance in the state space is studied as a measure of the dimensionality of a neural system[1,10,15–20]. Regardless of species and brain area, the dimensionality is drastically lower than the total number of recorded neurons (i.e., state-space dimension)[1], suggesting robust encoding of low-dimensional variables. Stringer et al.[10] showed that the dimensionality of visual cortical activity in mice can vary dynamically to encode precise visual input, seen as changes in the power law exponent of the explained variance. Such dynamical changes in dimensionality have not yet been demonstrated in other species.

Whether a subject has its eyes open or closed is known to affect the activity in the visual cortex, even in darkness[21–25]. In particular, the spectral power in the alpha frequency band (roughly 8–12 Hz) is known to decrease when the eyes are open, commonly known as alpha blocking[26–28]. Alpha blocking is usually attributed to desynchronisation[27] or oscillatory damping[28] within V1. However, the concrete pathway(s) triggering these phenomena, and the relation between eye closure and neural manifolds in V1, are still unknown.

The primary visual cortex (V1) is known to represent fine details of visual input at both single-neuron and population levels[12,29]. The visual system is hierarchical in nature, with information travelling from lower to higher areas (bottom-up) and vice versa (top-down), within specific frequency bands[30–32]. Top-down signals from V4 to V1 are known to mediate visual attention for figure-ground segregation and contour integration in macaque[33–36]. Recent evidence suggests that top-down signals can modulate neural manifold geometry and their dimensionality[37,38]. Naumann et al.[37] show *in silico* that top-down signals can rotate neural manifolds to maintain context-invariant representations. Dahmen et al.[38] show that recurrent connectivity motifs modulate the dimensionality of the cortical activity. As effective connectivity is input-dependent[39], a change in top-down input may therefore affect the dimensionality of neural activity. However, whether top-down signals modulate neural manifold geometry and dimensionality *in vivo* remains to be shown.

Here, we study the state space of the primary visual cortex of the macaque (N=3) during the resting state and its relation to the top-down signals from higher visual areas (V4, DP). We find that the population activity of macaque V1 is organised as two distinct high-dimensional neural manifolds, which are correlated with the behaviour (eye closure) of the macaques, but not related to any visual stimuli. The dimensionality of each of these manifolds is significantly different, with higher dimensionality found during the eyes-open periods than the eyes-closed periods. In addition, we estimated input from higher cortical areas to V1 and found that these top-down signals (in the form of LFP beta-band spectral Granger causality) are significantly stronger during the eyes-open periods, suggesting they play a role in modulating the neural manifolds and their dimensionality. Finally, we simulate a spiking neuron model under resting-state conditions and show that top-down signals can induce multiple manifolds by changing the firing modes of the network. Taken together, the data analysis and simulations show that top-down signals could actively modulate the V1 population activity, leading to two distinct neural manifolds of macaque visual cortical activity.

Results

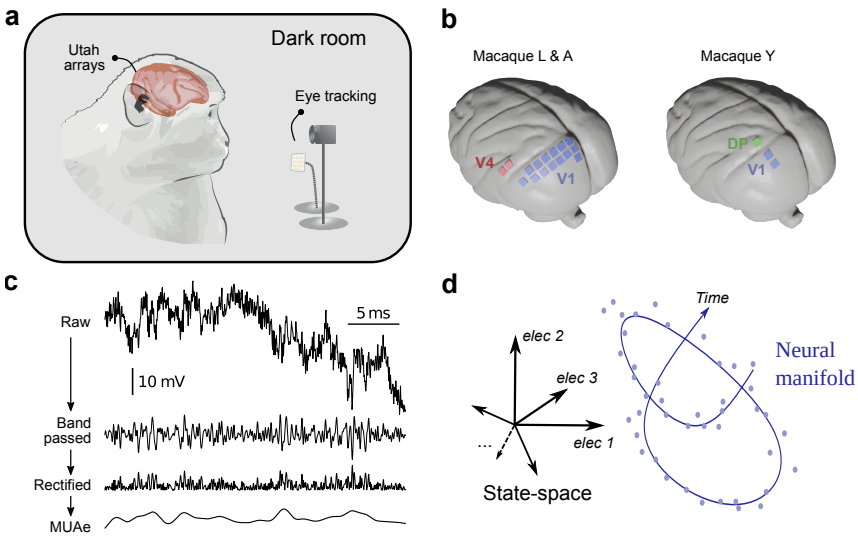

Figure 1: Overview of the experiment and neural manifold construction. **a** Illustration of the experimental setup. **b** Approximate locations of array implants in both experiments. Exact placement of the arrays differs slightly between subjects L and A. **c** Steps for obtaining the multi-unit activity envelope (MUAe)[40], used in this study. Band-pass filtering is performed between 500 Hz and 9 kHz, and the rectified signal is low-passed at 200 Hz to obtain the MUAe. **d** Schematic representation of state space and a neural manifold. Note that time is implicit within the neural manifold.

To explore the activity in the visual cortex, the intracortical electrical potential from the visual cortex of three rhesus macaques (*Macaca mulatta*) was recorded. The experiments simultaneously recorded the activity from V1 and V4 (macaques L & A)[41] and from V1 and DP (macaque Y, see Figure 1b)[42]. The recordings were made in the resting state, i.e., the macaques sat head-fixed in a dark room and were not instructed to perform any particular task. In this state the macaques often showed signs of sleepiness and kept their eyes closed for periods of variable duration. The right eye—contralateral to the site of neural recording—was tracked using an infrared camera, allowing the identification of periods of open or closed eyes. See methods Electrophysiological data from macaques L & A and Electrophysiological data from macaque Y for further details on the data acquisition and processing. The experimental setup and data processing steps are illustrated in Figure 1.

Two distinct neural manifolds in V1 correlated with eye closure

To explore the activity of the visual cortex, we characterise the high-dimensional population activity (between 64 and 800 electrodes, see Table 1 for details) for each area and macaque in terms of the downsampled (1 Hz) multi-unit activity envelope (MUAe)[40] (Figure 2a). We projected the population activity into a 3D space for visualisation using principal component analysis (PCA) (Figure 2b-d).

In V1, at least two distinct neural manifolds are apparent in the 3D projection space; sample session in Figure 2b-d, see Figure S1-S6 for all other sessions and subjects. We labelled the manifolds according to the sign of the log odds of a two-component Gaussian mixture model (see methods Neural manifolds and clustering and Outlier removal). The log odds represent the probability for a given data point to correspond to one manifold or the other. For the remainder of this paper we consider points with positive log odds to belong to 'manifold 0' (colour-coded in blue throughout the paper) and points with negative log odds to belong to 'manifold 1' (colour-coded in red throughout the paper).

To confirm that the two manifolds in the lower-dimensional projection are not an artefact of the dimensionality reduction, we estimated the Betti numbers of the high-dimensional population activity using persistent homology (see Methods, Topological data analysis). The persistence barcodes show that at least two independent generators of the H_0 homology groups exist in the high-dimensional population activity, corresponding to two connected components (Figure S7), i.e., two distinct neural

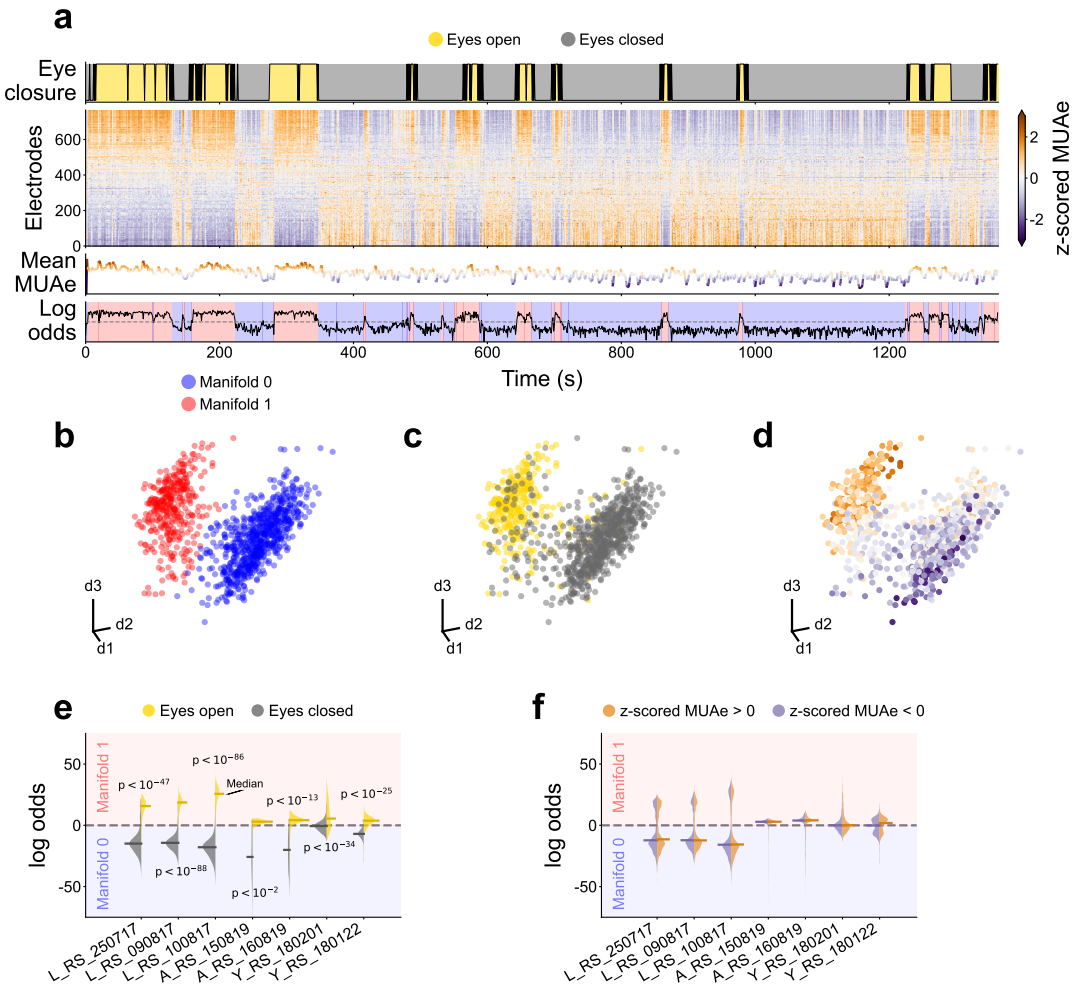

Figure 2: Two distinct neural manifolds in V1 correlated with eye closure. **a** Overview of the experimental data from session L_RS_250717. From top to bottom: Time evolution of the eye signal; the z-scored MUAe signal for each electrode (electrodes ordered by their correlation with the eye signal); the mean z-scored MUAe at each time point; and the log odds overlaid with the most likely manifold (two clusters, Gaussian mixture model). **b, c, d** First three principal components of the MUAe population activity. Colours indicate the manifold identified via the log odds of a Gaussian mixture (**b**), the eye closure (**c**) and the mean z-scored MUAe (**d**). Each dot represents a different point in time. Outliers were excluded from the neural manifolds shown in **b**–**d**, see Outlier removal. **e, f** Violin plots of the distribution of the log odds across epochs, respectively distinguished according to eye closure (**e**, result of a logistic regression test shown) and z-scored MUAe (**f**). Horizontal bars indicate medians of the distributions.

manifolds. Thus, we confirm that the two manifolds observed in the 3D projection are inherent to the high-dimensional space.

Additionally, we tested whether the observed manifolds could be an artefact of the MUAe signal. We spike-sorted one session (L_RS_250717) with a semi-automatic method and analysed the population activity resulting from the single-neuron firing rates (Figure S8). The spiking activity also displayed two manifolds, in agreement with the MUAe signals.

While the activity of visual cortex is mainly driven by visual input, whether and to what extent it is separately modulated by eye closure is unclear. Marking data points on the V1 manifolds with the eye closure signal (Figure 2b) reveals that one manifold strongly relates to the eyes-open periods, whereas the other manifold strongly relates to the eyes-closed periods.

To confirm the correlation between eye closure and manifolds, we tested the differences between the eyes-open and eyes-closed periods using a twofold approach. First, we performed a logistic regression between the eye closure signal and the log odds, revealing a significantly higher than chance correlation

in all sessions (Figure 2e). Second, we visualised the distribution of the log odds during the eyes-open and eyes-closed periods separately; showing a clear correspondence between the eye closure and the sign of the log odds in most cases (Figure 2e). Taken together, the logistic regression and the log-odds distributions demonstrate that membership of a point in state space in one of the two V1 manifolds is closely related to eye closure. Given this, we will refer to the manifolds as the eyes-open manifold or the eyes-closed manifold.

The existence of two separate manifolds could be trivially explained if the MUAe activity levels were significantly different between the eyes-open and eyes-closed periods, and the manifolds simply reflected the population activity level. To examine this possibility, we checked whether higher-activity epochs uniquely correspond to one of the manifolds. The violin plots of the full data distribution—based on the z-scored MUAe shown in Figure 2a—show that there is no clear separation into two manifolds (Figure 2f). Additionally, we visualised the 2D histograms of z-scored MUAe against log odds (Figure S9). Both the violin plots and the 2D histograms suggest that the activity level alone does not fully explain the presence of the two neural manifolds in macaque V1.

For completeness, we also visualised the population activity from V4 and DP (Figure S10, S11). In contrast to V1, the population activity in areas V4 and DP does not appear to contain two distinct neural manifolds. We also tested the relationship between neural activity and eye closure in V4 and DP (Figure S12), using the same procedure as for V1. Although some correlation is observed between eye closure and log odds, the violins reveal no clear manifold separation. Thus, we conclude that the observed manifolds are restricted to V1 and are not present in V4 or DP.

Higher dimensionality during eyes-open periods, primarily due to decorrelation

To further understand the functional role and implications of the observed neural manifolds in V1, we studied their dimensionality. We used the participation ratio (PR, Equation 1), which is defined via the percentage of variance explained by the principal components of the covariance matrix[19,38]. The PR can be rewritten in terms of the statistics of the covariance matrix

$$\mathrm{PR} = \frac{\left(\sum_i \lambda_i\right)^2}{\sum_i \lambda_i^2} = \frac{N}{1 + v^2 + (N-1)(m^2 + s^2)} \tag{1}$$

where λ_i are the eigenvalues of the covariance matrix and N is the number of electrodes. v, m, and s are the ratios between the standard deviation of auto-covariances, average cross-covariances, and the standard deviation of cross-covariances with respect to the average auto-covariances, respectively. See Dimensionality for detailed methods.

To study the dimensionality, we computed the time-varying PR from the z-scored MUAe signals, by calculating the PR for windows of 30 s width (1 s steps, thus 29 s of overlap with adjacent windows), see Figure 3a. Stronger MUAe activity is typically associated with higher variance, which may bias the results toward higher dimensionality. We avoided bias due to the varying activity level by normalising the data (via z-scoring) within each window. We found that there is a strong correlation between the log odds and the time-varying PR (Figure 3b) and compared the PR values between the two manifolds using a Mann-Whitney U test (Figure 3c). The correlation and tests show that the dimensionality is higher during the eyes-open periods, consistently for all data sets.

To further support this finding, we also show the distribution of the variance explained by each of the principal components (PC) of the MUAe data, depicted on a log-log scale in Figure 3d. We fitted a power law to the PC variances and report the exponent α (Figure 3e). A higher α indicates faster decay of the curve, i.e., lower dimensionality. The power law exponents are in agreement with our sliding window approach: We observe higher dimensionality during eyes-open than during eyes-closed periods (Figure 3e).

To narrow down the reason causing the dimensionality changes, we computed v^2, $(N-1)m^2$, and $(N-1)s^2$ and observed that the changes in $(N-1)m^2$—i.e. the average cross-covariances—dominate the PR differences between the eyes-open and eyes-closed periods (Figure 3f). Thus, the main reason for the observed dimensionality changes is decorrelation during the eyes-open periods.

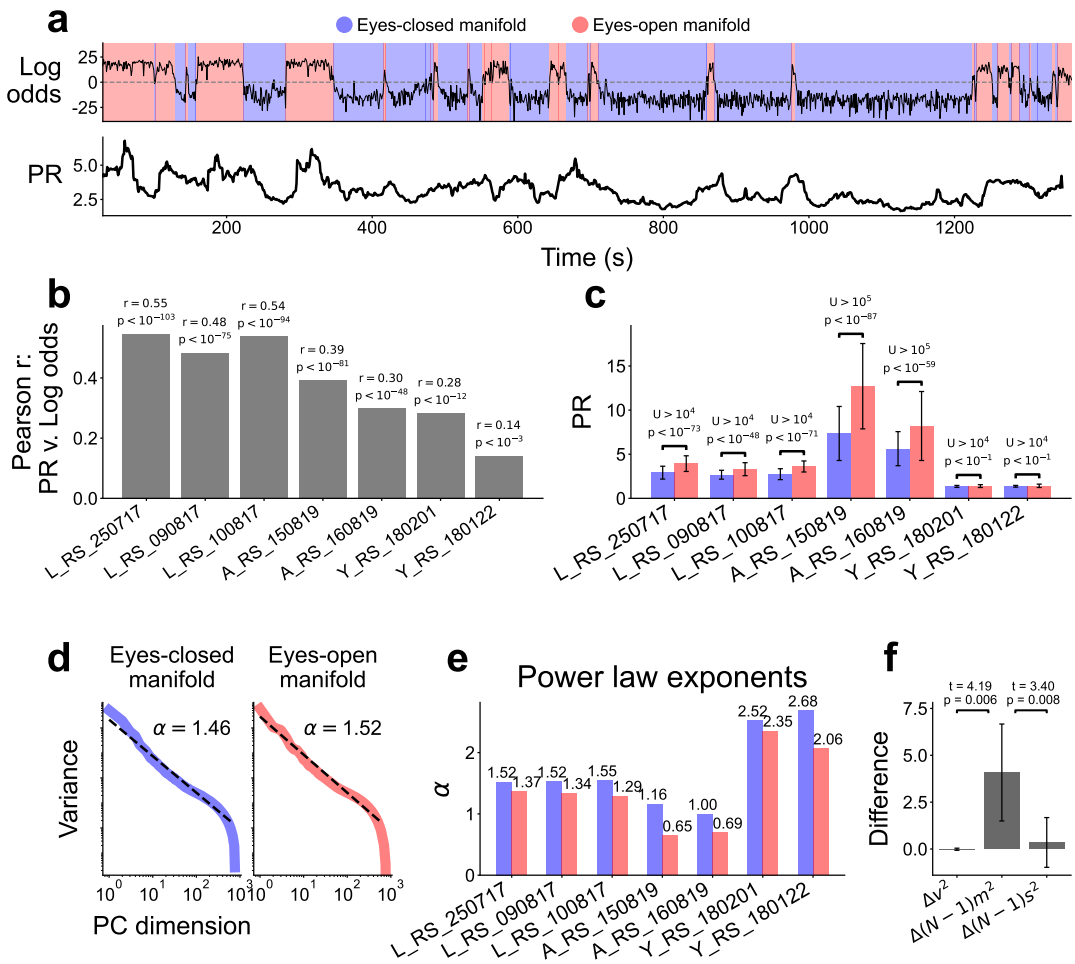

Figure 3: Higher dimensionality during eyes-open periods. **a** Log odds and participation ratio (PR) for session L_RS_250717. The PR was calculated on a sliding window of 30 s width. **b** Pearson correlation between log odds and PR. **c** Comparison of PR between neural manifolds (Mann-Whitney U test). **d** Distribution of principal components and their explained variance on a log-log scale, respectively for the entire session (L_RS_250717) and for each manifold. Power law exponent α estimated over the ranges where the curves approximate a power law. **e** Comparison of power law exponents for the two neural manifolds in all sessions. The eyes-open manifold always had a smaller exponent, indicating a higher dimensionality. **f** Differences in the terms of the PR function between eyes-open and eyes-closed, results of Welch's t-test across sessions shown. The quantities are related to the standard deviation of auto-covariances (v^2), average cross-covariances ($(N-1)m^2$), and standard deviation of cross-covariances ($(N-1)s^2$)

Top-down signals from V4 to V1 are present in the form of beta-band spectral Granger causality

In search of an internal mechanism that may modulate the neural manifolds and their dimensionality, we turned our attention to cortico-cortical interactions. Since signatures of top-down activity have previously been reported in the beta frequency band (roughly 12–30 Hz)[31,44], we use spectral Granger causality to measure top-down signals.

To determine whether top-down signals are present in our data, we calculated the coherence and Granger causality between every pair of V1-V4 and V1-DP electrodes (see Coherence and Granger causality)—using the local field potential (LFP). Figure 4a,b show the coherence and Granger causalities for a sample pair of electrodes. To quantify the cortico-cortical signals, we searched for peaks in the coherence and Granger causality, using an automatic method (see Methods, Peak detection). We detected beta-frequency Granger causality peaks in around 0.5% of all V1-V4 electrode pairs, predominantly in the top-down direction (Figure S13). We only found beta-band bottom-up interactions

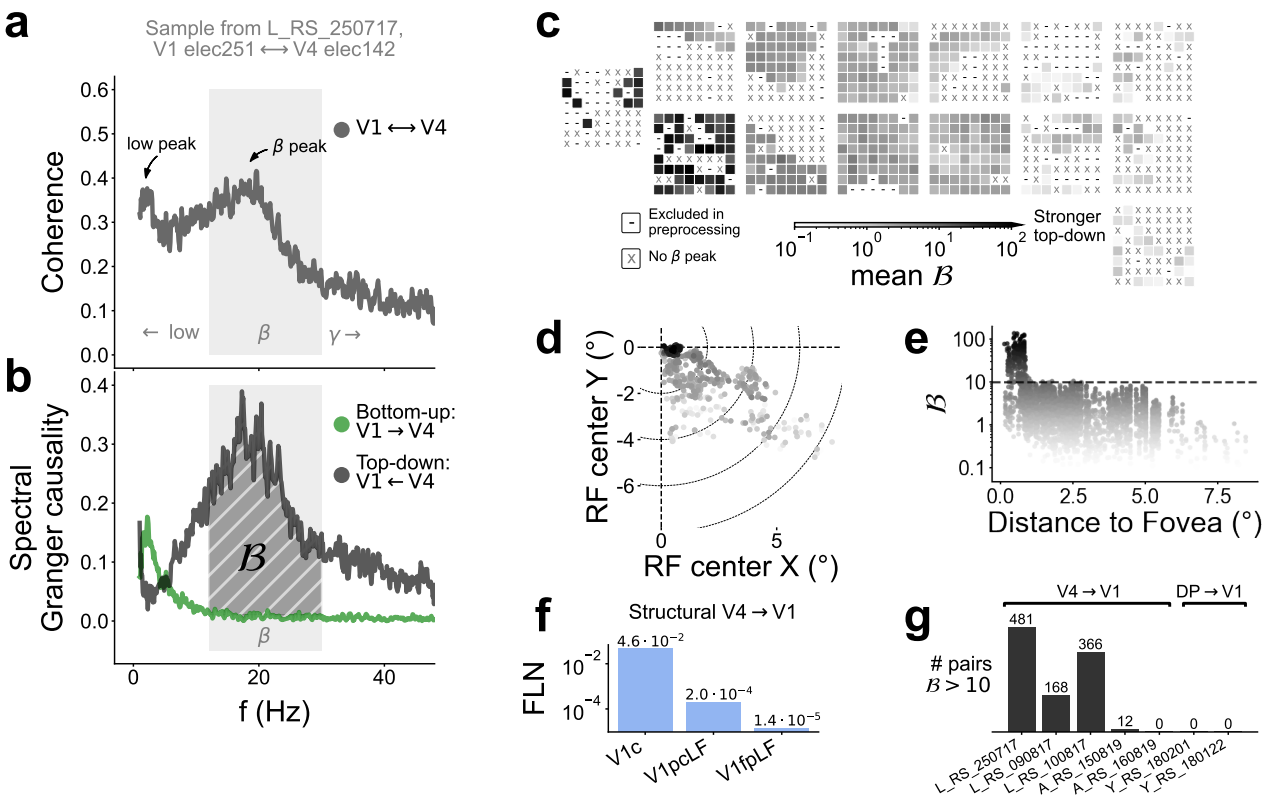

Figure 4: Inter-area coherence and spectral Granger causality. **a** Representative sample of coherence between V1 and V4 (electrodes 242 and 142, respectively). Low-frequency and beta-band peaks indicated. **b** Representative sample of spectral Granger causality. **c** Schematic representation of the electrode locations overlaid with the mean top-down signal strength $\mathcal{B}$ per electrode (see Coherence and Granger causality for a description of $\mathcal{B}$). **d** Receptive field (RF) map overlaid with the mean $\mathcal{B}$ per electrode. Stronger $\mathcal{B}$ is found around the foveal region of V1. **e** Mean $\mathcal{B}$ displayed against the distance from the fovea. **f** Log of fraction of labelled neurons (FLN) from V4 to V1 (data from tract-tracing experiments[43]). V1 subdivisions represent c: central (foveal region), LF: lower visual field, pc: peri-central, and fp: far periphery. The strongest connectivity exists from V4 to V1c, in agreement with our measurements. **g** Number of electrode pairs with strong ($\mathcal{B} > 10$) top-down signals detected in each session.

in V1-DP electrode pairs.

For the electrodes with a beta causality peak, we estimated the causality strength $\mathcal{B}$ (Equation 5). The electrodes with their receptive field (RF) closer to the fovea show substantially higher $\mathcal{B}$ (Figure 4c–e, Figure S14), in agreement with a previous structural connectivity report[43] (Figure 4f). To disregard potential spurious Granger causality peaks, we restrict all further analysis of the top-down signals to the strongest interactions, by setting a threshold of $\mathcal{B} > 10$ (Figure 4g). We found no bottom-up V1-to-V4 signals with high strength in the beta frequency band.

Thus, we found top-down signals from V4 to V1, in agreement with previous studies[31,44]; but we did not find strong signals from DP to V1 in our data. V4-to-V1 signals are therefore strong candidates for the modulation of the neural manifolds and their dimensionality.

Stronger top-down signals from V4 to V1 during eyes-open periods

To elucidate the behavioural relevance of the V4-to-V1 top-down signals, we examined how the spectral power, coherence, and Granger causality change of the LFP in relation to eye closure.

We extracted the LFP data for each behavioural condition and concatenated the data within the same condition. This approach could introduce some artefacts, which we expect to be minor in view of the very small number of transitions in comparison with the number of data samples (500 Hz resolution). Both in V4 and V1, we find that the spectral power at low frequencies (< 12 Hz) is higher

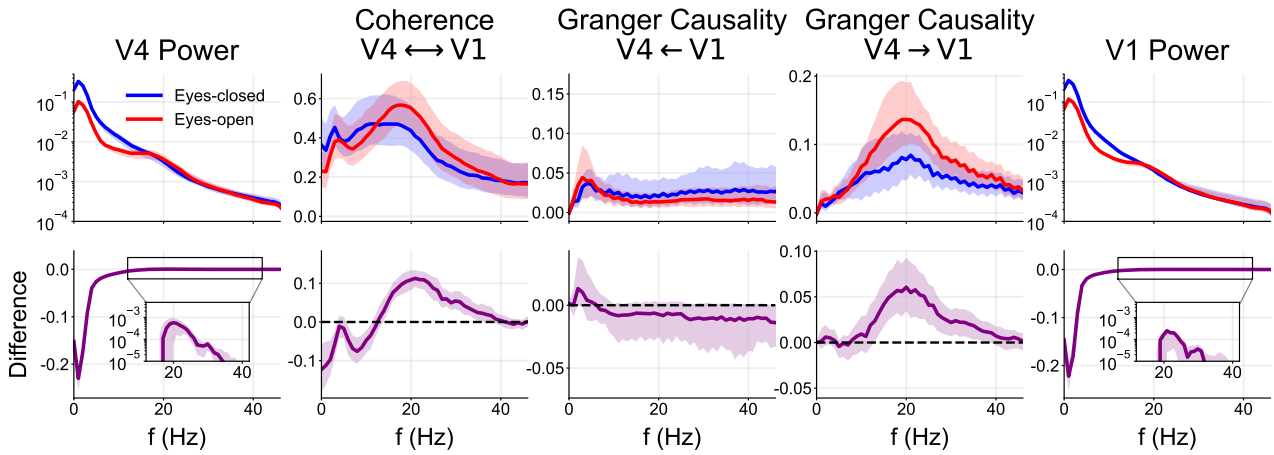

Figure 5: Spectral power, coherence, and Granger causality of the LFP for the electrodes with high causality strength ($\mathcal{B} > 10$) in session L_RS_250717, see Figure S15 for all other sessions. The data for each behavioural condition (eyes-open/closed) were concatenated and their metrics reported separately. Thick line shows median across electrodes (or pairs of electrodes) and shading indicates the 25th to 75th percentile (top row). The difference between eyes-open and eyes-closed periods was calculated for each electrode or pair of electrodes (bottom row).

during the eyes-closed periods, whereas the power in the beta band (12–30 Hz) is slightly higher during eyes-open periods (Figure 5, Figure S15). Spectrograms of the V1 LFP power confirm the reduction in low-frequency power during eyes-open periods (Figure S17). The coherence in the beta band is higher during the eyes-open periods, with the peak shifted to higher frequencies compared to the eyes-closed condition. Notably, the top-down Granger causality is substantially higher in the beta band during the eyes-open periods.

In order to confirm our observations, we also computed spectrograms of the Granger causality using a 10-second sliding window (Figure S16a). Statistical tests (Welch's t-test) of the difference between bottom-up and top-down Granger causality, ΔGC, confirmed a shift toward top-down interactions during the eyes-open periods compared to the eyes-closed periods, for a vast majority of all electrode pairs (Figure S16b,c). Thus, we found higher beta-band Granger causality during eyes-open periods using two different approaches.

We further tested whether the top-down signals were correlated with gaze direction and eye movements (Figure S18), to rule out the presence of any visual stimuli—despite the experiments being performed in a dark room. No clear trend could be observed, thus indicating no relation between gaze direction and top-down signals. This finding suggests that the visual scene is not the source of the observed top-down signals.

In conclusion, the time-dependent spectral analysis reveals large variations of power and Granger causality. On the one hand, the spectral power at low frequencies decreases during eyes-open periods, consistent with the well-known alpha blocking phenomenon[26–28]. On the other hand, the V4-to-V1 top-down signals are strongest during the eyes-open periods. The time-varying top-down beta causality strength did not substantially correlate with gaze direction or eye movements, suggesting no relation between the top-down signals and the visual scene; as expected in a dark room. Taken together, these results suggest that V4-to-V1 signals modulate V1 activity, contributing to a different state-space manifold with increased dimensionality.

Discussion

In this paper, we presented three novel findings in the primary visual cortex (V1) of macaques during the resting state: two separate manifolds in the state space associated with eye closure (Figure 2); higher dimensionality due to lower mean cross-correlations during eyes-open periods (Figure 3); and the presence of stronger top-down signals from V4 to V1 during the eyes-open periods, primarily targeting

the foveal region of V1 (Figure 4, Figure 5). In addition, we observed lower power at frequencies below 12 Hz during the eyes-open periods (Figure 5, Figure S17), consistent with the well-known alpha blocking effect[26].

We observed that two distinct manifolds appear in the state space of macaque V1—but not V4 nor DP—during the resting state for all subjects and sessions, for both MUAe and spike data (Figure 2, S1-S6, Figure S8), and are correlated with eye closure (Figure 2e,f). The manifolds were not just an artefact of the three-dimensional projection used for visualisation, as we confirmed they also exist in higher dimensions with persistent homology (Figure S7). Previous work in mice has shown that the visual cortex represents a myriad of behaviours in the resting state, such as facial movements or running[45]. However, a similar study on the macaque showed that the macaque visual cortex is very specific to vision, and minimally driven by spontaneous movements[46]. Thus, we do not expect the neural manifolds of V1 to be strongly affected by any behaviour other than visual behaviour, in agreement with our finding that eye closure neatly explains the two manifolds.

Our findings could in principle explained by the presence of complex visual stimuli that would alter the population dynamics and cortico-cortical communication. However, we are certain that no strong visual stimuli are present in the visual field, due to the very dark environment of the recording room. Additionally, we performed several analyses to control for activity levels (Figure 2 and Figure S9) and gaze direction Figure S18. Furthermore, the original data for Macaques L and A includes an extensive evaluation of data quality, which excluded all electrodes that did not strongly respond to visual stimuli[41]. Thus, all the electrodes included in our analysis (from Macaques L and A) would strongly respond if there were strong visual stimuli, but we observed no such responses in the MUAe activity (see Figure 2). We are therefore certain that the visual input is faint or nonexistent, which implies that the observed neural manifolds must be induced by some other internal mechanism.

Further characterisation of the activity in the different manifolds revealed that the neural dimensionality is manifold-dependent (Figure 3). We observed higher dimensionality in the eyes-open manifolds across all macaque and sessions. Our measured dimensionality is in agreement with previous reports on the visual cortex[10,16]. Previous work has also shown higher dimensionality in the primary motor cortex during eyes-open than eyes-closed periods[47], analogous to our findings in the visual cortex.

We hypothesised that top-down signals from higher cortical areas could be the modulatory mechanism responsible for the changes observed in the neural manifold and dimensionality of V1 activity. Indeed, we found that there are strong top-down signals from V4 to V1 (Figure 4), targeting particularly the foveal region of V1, in agreement with structural connectivity[43]. We also found the top-down signals to vary over time, with increased presence during the eyes-open periods (Figure 5). In agreement with our findings, previous studies found that cortico-cortical top-down signals between V1 and V4 are predominantly present in the beta (12–30 Hz) frequency band, while bottom-up signals between V1 and V4 are present in the delta/theta (< 8 Hz) and gamma (> 30 Hz) bands[31,44]. In our analysis we did not find any gamma band causality (Figure 4), likely because our recordings were from the deep cortical layers (in macaque L and A the electrodes were 1.5 mm long, putatively recording mostly from layer 5) and gamma oscillations are known to be weak in layer 5 of the visual cortex[30,48]. In contrast to our findings, van Kerkoerle et al.[30] reported that top-down signals appear in the alpha (8–12 Hz) frequency range. Whether the specific top-down and bottom-up frequencies generalise to the whole cortex is unclear. Instead, Vezoli et al.[44] postulate overlapping modules of certain frequencies (alpha, low-beta, high-beta, and gamma) that differ across cortical areas. Our findings are consistent with the work by Semedo et al.[49], who suggested that bottom-up signals dominate during visual stimulation and top-down signals dominate in the absence of visual stimuli—note that in their work the eyes were always open. We did not find top-down signals from DP to V1, possibly due to the electrodes used in macaque Y being 1 mm long, thus likely recording from layer 4. In fact, top-down connections do not originate in nor target layer 4 of cortex[50].

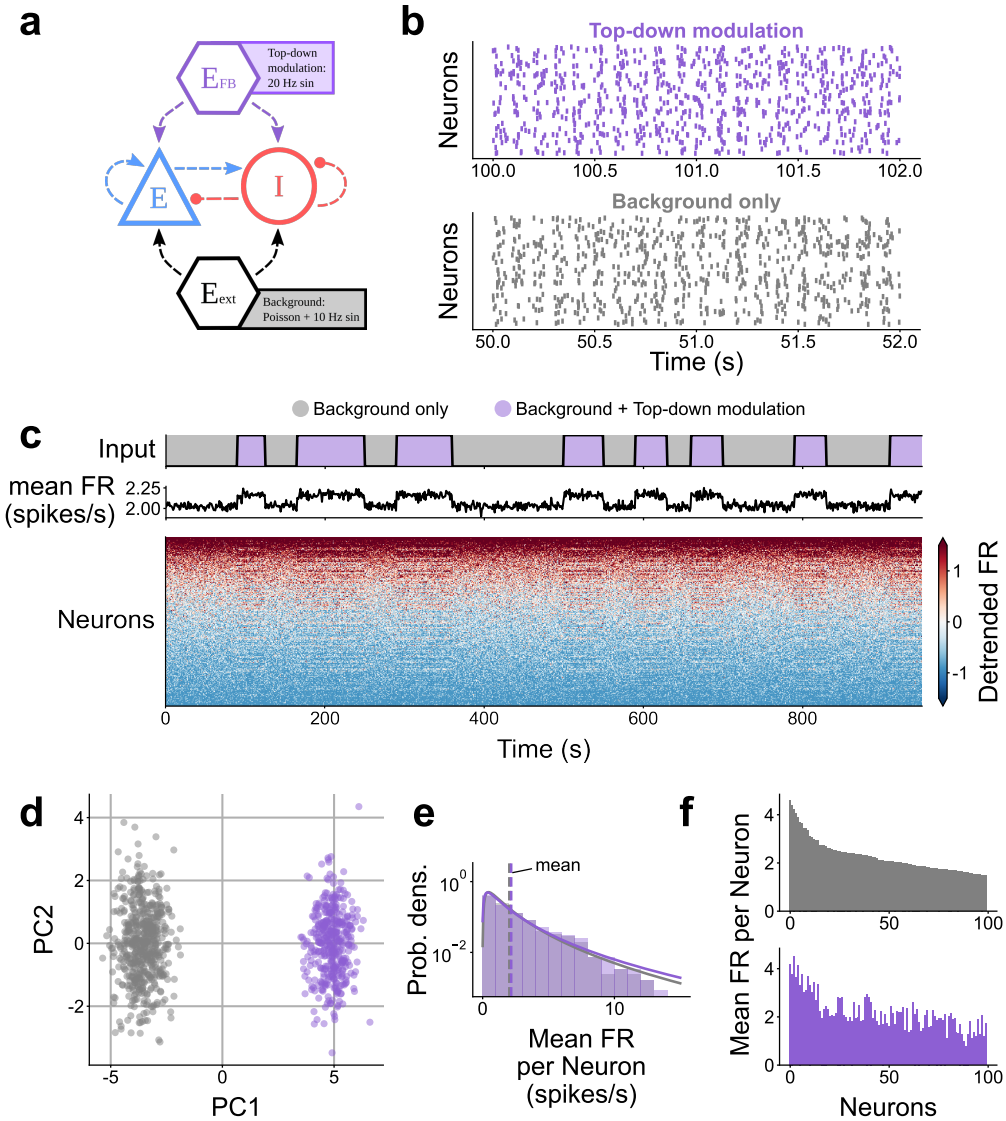

Figure 6: Simulation of a balanced spiking neural network with top-down modulation. **a** Diagram of balanced random spiking neural network. Background input is provided constantly, top-down signals are provided intermittently. **b** Sample raster plots show spiking activity in the different input regimes. **c** Time evolution of input regimes and mean firing rate (FR). **d** First two principal components of the firing rate (binsize = 1 s). Colours indicate the different input regimes. **e** Distribution of mean firing rate per neuron is almost identical between the two regimes. **f** Mean firing rate of the 100 most active neurons. The top-down modulation changes the mean firing rates of each neuron, in both the positive and negative directions, leading to the observed distinct manifolds.

In the present study, it was not possible to test directly from the experimental data whether the V4-to-V1 signals are responsible for the modulation of V1 dynamics. Future studies could perform such a test by a reversible inactivation of the V4-to-V1 pathway, such as via reducing the temperature of V4[51,52], injecting a GABA agonist (e.g., muscimol, bicuculline)[53–55] or using targeted optogenetic suppression[56]. These techniques have been successfully applied to study the suppression of cortico-cortical communication; however, to best of our knowledge, they have not been used to study the effects of macaque V4-to-V1 signals in the resting state.

Numerical simulations offer an alternative approach to study the effect of top-down signals in spiking neural networks. We thus performed preliminary simulations of a simple spiking neuron model—of the well-known Brunel type[57]—to ascertain whether V4-to-V1 signals can modulate the neural manifolds (Figure 6). Modelled top-down signals, in the form of sinusoidal oscillating inhomogeneous Poisson processes, led to a different neural manifold in the network activity when a subset of the network

neurons was targeted (Figure 6d). These changes were not due to the increase in firing rate caused by the additional top-down input, but rather due to the activation of different neuron patterns in the model (Figure 6e,f). We limited the analysis of the model to the presence of neural manifolds, because our model was ill-suited to study the dimensionality, given that average cross-correlation is known to cancel out in balanced EI networks[58,59]. Future work could use more complex models—such as clustered networks[58,60,61]—to study the effects of correlated inputs with realistic power spectra on the dimensionality and elucidate whether the top-down signals can directly induce the observed increase in the dimensionality during the eyes-open periods.

Taken together, our data analysis and simulations suggest that top-down modulation alone is sufficient to cause the distinct neural manifolds in V1 activity. Nevertheless, sustaining the different V1 manifolds might involve additional mechanisms, such as neuromodulation or adaptation of recurrent connectivity via short-term plasticity. Previous work suggests that N-methyl-D-aspartate (NMDA) receptors are central to the top-down communication from V4 to V1[34,62]. Interestingly, targeted pharmacological deactivation of NMDA receptors in macaque V1 leads to the suppression of alpha blocking[28] and absence of decorrelation during eyes-open[63]; both of which are correlated with the increased V4-to-V1 signals in our data. In addition, the top-down signals are not constant throughout the eyes-open periods (Figure S16), but the slow timescale of the NMDA receptors could help sustain the eyes-open manifold, even if the top-down input fades. Thus, we speculate that the top-down connections preferentially target the NMDA receptors in V1 neurons, leading to the observed alpha blocking and decorrelation. Additionally, recurrent connectivity, in the form of cell-type-specific motifs, has also been shown to constrain and control the dimensionality of brain networks[38], which could emerge in the effective connectivity of the network as a result of the top-down input.

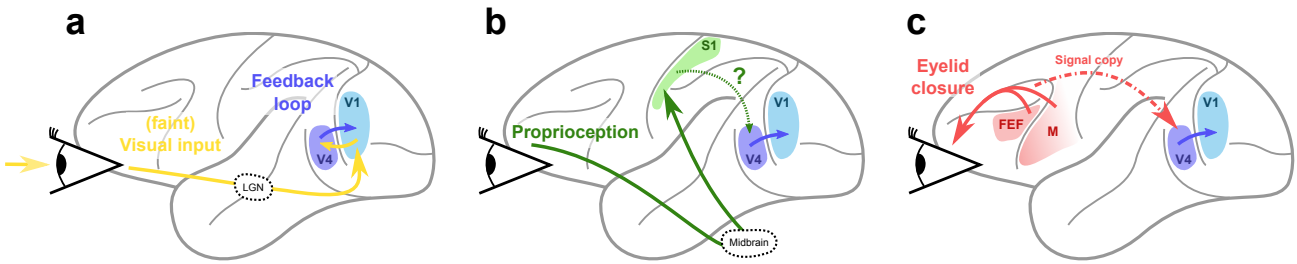

Figure 7: Proposed communication pathways for V1 modulation via V4. **a** Visual input directly to V1 triggering a cortico-cortical feedback loop. **b** Proprioception of eyelid muscles via the midbrain (possibly superior colliculus) and somatosensory cortex. **c** Cortico-cortical communication of motor commands.

If the V4-to-V1 signals convey behavioural information, then how does such behavioural information reach V4 in the first place? We explore the possible communication pathways that lead to the observed V4-to-V1 signals, illustrated in Figure 7. We identified three main candidates: the visual stimulus (or absence thereof) from the retina to V1; the proprioception of eyelid muscles via the somatosensory cortex; and the voluntary motor commands for eye closure. The first proposed pathway involves visual stimuli being transmitted from the retina to V1 via the lateral geniculate nucleus (LGN). The absence of stimuli could be the reason for the observed changes in the V1 activity, whereas the presence of visual stimuli could trigger a V1-V4 feedback loop. However, the macaques in our experiments had very little to no visual input, even during eyes-open periods, since they were sitting in a dark room. Additionally, we found no consistent difference in MUAe activity levels between the eyes-open and eyes-closed manifolds (Figure 2). The second proposed pathway involves proprioception of the eyelid that could inform the cortex when the eyes are closed and trigger the activity changes in V1. Mechanoreceptors in the eyelid activate the oculomotor nerve projecting to the midbrain (possibly to the superior colliculus)[64], eventually entering the cortex via the somatosensory area (S1)[65]. From S1 the signal could find its way to V1 via several cortico-cortical pathways, potentially including neurons in V4; however this mechanism might be relatively slow, given the absence of direct connections from S1 to V1 or V4[66]. Furthermore, the shortest known S1-to-V1 cortico-cortical pathway does not involve V4, rendering this pathway a rather weak candidate. The third and final proposed pathway involves

voluntary eyelid closure which is initiated by the ventral motor cortex and the frontal eye field (FEF). The eyelid closure and eye movements may be communicated to the visual cortex via cortico-cortical connections or the superior colliculus. Given that V4 is part of the fronto-parietal network (with strong FEF⇄V4 connections)[66,67], the eye movement signals could easily reach V4, which could then modulate the V1 activity.

The hypotheses from Figure 7 are not necessarily mutually exclusive, and could all play a role in the modulation of V1 activity. To understand which pathways are most relevant to sustain the manifolds, we had a closer look around the manifold transitions (Figure S19) by looking at the MUAe signals at a high temporal resolution (1 kHz). For the eye-opening transitions, we observed that sometimes V1 MUAe activity precedes V4 activity, in agreement with the feedback loop hypothesis (Figure 7a); whereas in other cases V4 activity precedes V1 activity, in agreement with the hypotheses shown in Figure 7b,c. In the eye-closing transitions, the activity from V1 and V4 appeared to be simultaneous. The number of transitions was relatively small, which did not allow for a quantitative analysis of the transitions between the two manifolds. Further work could revisit this issue by looking at longer recordings including larger numbers of transitions between eyes-open and eyes-closed periods.

Given the complex mechanisms that seem to be involved in ensuring that V1 population activity adjusts to eye closure, it seems likely that it has a functional benefit. First of all, if the eyes are closed, no visual stimuli are processed and V1 firing rates are reduced to save energy. On the other hand, when the eyes are open, higher-dimensional activity might be advantageous for better encoding visual stimuli, which are know to have a high dimensionality[10]. This could thus facilitate visual processing. Previous work showed that spectral power in the alpha band (8–12 Hz) is inversely correlated with visual recognition performance in human subjects[27,68]: lower alpha power was associated with better performance in a visual discrimination task. Our results suggest that the change in neural manifolds and dimensionality are directly correlated with the decrease in alpha power (Figure 5, Figure S17). Future work could study the relation between the dimensionality, alpha power, and visual performance (e.g., response latency to different images) to determine the functional relevance of our findings.

In conclusion, we provide *in vivo* evidence for the modulation of neural manifolds by cortico-cortical communication, which we hypothesise could enable more efficient responses to visual stimuli. Our analysis and previous results suggest that the eyes-open manifold—together with the corresponding dimensionality and spectral power changes—constitutes a visual stand-by mode, which is modulated by top-down input from V4 and other internal mechanisms.

Methods

Electrophysiological data from macaques L & A

We used publicly available[41] neural activity recorded from the neocortex of rhesus macaques (N=2) during rest and a visual task. The macaques were implanted with 16 Utah arrays (Blackrock microsystems), two of them in visual area V4 and the rest in the primary visual cortex (V1), with a total of 1024 electrodes. The electrodes were 1.5 mm long, thus recording from the deeper layers, likely layer 5. The recording system recorded the electric potential at each electrode with a sampling rate of 30 kHz. A full description of the experimental setup and the data collection and preprocessing has already been published[41]; here we only provide the details relevant to this study.

Three resting-state (RS) sessions were recorded per macaque, during which the subjects did not have to perform any particular task and sat in a quiet dark room. Pupil position and diameter data were collected using an infrared camera in order to determine the direction of gaze and eye closure of the macaques. On the same days as the RS recordings, a visual response task was also performed. The visual response data were used to calculate the signal-to-noise ratio (SNR) of each electrode, and all electrodes with an SNR lower than 2 were excluded from further analysis. Additionally, we excluded up to 100 electrodes that contributed to high-frequency cross-talk in each session, as reported in the original data publication. The sessions, duration and number of electrodes per subject are listed in Table 1.

The raw neural data were processed into the multi-unit activity envelope (MUAe) signal and local field potential (LFP). To obtain MUAe data, the raw data were high-pass filtered at 500 Hz, rectified, low-pass filtered at 200 Hz, and downsampled to 1 kHz. Finally, the 50, 100, and 150 Hz components were removed with a band-stop filter in order to remove the European electric grid noise and its main harmonics. During the recordings, the macaque's head was held in position with a custom-made headstage. To obtain the LFP data, the raw data was low-pass filtered at 250 Hz, downsampled to 500 Hz and a band-stop filter was applied to remove the European electric grid noise (50, 100, and 150 Hz).

The MUAe and LFP data for each array were already provided by the original authors in the open-source .nix format, which uses python-neo data structures to hierarchically organise and annotate electrophysiological data and metadata. The metadata, such as the cross-talk removal or the positions of the arrays in the cortex, were provided in the .odml machine- and human-readable format, which were incorporated into the python analysis scripts.

Spike sorting

The raw data from one session (L_RS_250717) were spike-sorted using a semi-automatic workflow with Spyking Circus—a free, open-source, spike-sorting software written entirely in Python[69]. An extensive description of the methods of this algorithm can be found in their publication, as well as in the online documentation of Spyking Circus[1].

Roughly, Spyking Circus first applied a band-pass filter to the raw signals between 250 Hz and 5 kHz. Next, the median signal across all 128 channels that shared the same reference (2 Utah arrays) was calculated and subtracted, in order to reduce cross-talk and movement artefacts. The spike threshold was set conservatively, at eight times the standard deviation of each signal. After filtering and thresholding, the resulting multi-unit spike trains were whitened—removing the covariance from periods without spikes to reduce noise and spurious spatio-temporal correlations. After whitening, a subsample of all spike waveforms is selected, reduced to the first five principal components, and clustered into different groups with the k-medians method. Finally, all spikes in each electrode are assigned to one of the waveform clusters based on a template fitting algorithm, which can also resolve overlapping waveforms.

After the automatic sorting, the waveform clusters were manually merged and labelled as single-unit activity, multi-unit activity, or noise. Only single-unit activity (SUA) spike trains were included in this study. The waveform signal-to-noise ratio (wfSNR) was calculated for all SUA, and those with a wfSNR < 2 or electrode SNR < 2 (from the visual response task) were excluded from the analysis.

[1] spyking-circus.readthedocs.io

Table 1: Summary of subjects and recordings included in this study.

Subject	Session	Duration (s)	Areas	Electrodes	Clean Electrodes
L	L_RS_250717	1363	V1	896	765
			V4	128	116
L	L_RS_090817	1321	V1	896	761
			V4	128	116
L	L_RS_100817	1298	V1	896	774
			V4	128	118
A	A_RS_150819	2278	V1	896	402
			V4	128	11
A	A_RS_160819	2441	V1	896	369
			V4	128	9
Y	Y_RS_180122	906	V1	64	42
			DP	36	25
Y	Y_RS_180201	699	V1	64	44
			DP	36	24

Electrophysiological data from macaque Y

In addition to the published data from macaques L & A, we also used an unpublished data set from one additional rhesus macaque (N=1). Neural activity was recorded during rest and during a visuomotor integration task. The recording apparatus is described elsewhere[42]. The macaque was implanted with five Utah arrays (Blackrock microsystems), two of them in the primary visual cortex (V1), one in dorsal prelunate cortex (area DP), one in area 7A and one in the motor cortex (M1/PMd). In this study we only included the 6x6 electrode arrays from V1 (two arrays) and DP (one array), for a total of 108 electrodes. The electrodes were 1 mm long, thus recording from the central layers, likely layer 4. The recording system recorded the electric potential at each electrode with a sampling rate of 30 kHz.

Two resting-state (RS) sessions were recorded, during which the macaque did not have to perform any particular task and sat in a quiet dark room. Pupil position and diameter data were collected using an infrared camera in order to determine the gaze direction and eye closure of the macaque. See Table 1 for an overview of the sessions used in this study.

MUAe and LFP signals were computed using the same procedure as for the other data sets.

Neural manifolds and clustering

The MUAe data were downsampled to 1 Hz and arranged into a single array, with between 50 and 900 recording locations per session.

In order to visualise the data, we used a standard dimensionality reduction technique (principal component analysis, PCA) to reduce the neural manifold to 3D. The clusters observed in the RS sessions were labelled using a two-component Gaussian mixture model on the 3D projection. The clustering method provides the log odds, i.e., the chance that any given point belongs to one cluster or the other. The log odds captures the multi-cluster structure of the manifold in a single time series; thus, we consider it to be an identifier of the V1 manifolds.

Outlier removal

The neural manifolds in our analysis are a collection of time points scattered across the state space. In the data some time points appear very distant from all other points, which we associate with noise and we therefore seek to remove them. To identify the outliers we used a procedure similar to the one used by Chaudhuri et al.[5]. First, we calculated the distance matrix of all points to each other, and took the

1st percentile value from the distance distribution, D_1. We then estimated the number of neighbours that each point had within D_1 distance, and finally discarded the 20 percent of points with the fewest neighbours.

Topological data analysis

We used persistent homology to confirm that the lower-dimensional structures that we observed in the 3D projection of the neural manifolds are in fact topological features of the data and not just an artefact of the dimensionality reduction. Before computing the persistence barcodes we projected the data into a 10D subspace using the isomap technique[70]. The method aims at approximately preserving the geodesic distance between data points (that is the shortest path between two points on the neural manifold) and thus is suited for reducing the dimensionality of the data when applying a topological data analysis. The analysis on the 10D data showed qualitatively equivalent results to the full-dimensional data, while requiring a much shorter computation time.

To calculate the persistence barcodes of the Vietoris-Rips complex of the neural manifold we used an efficient open-source implementation (Ripser[2]). Briefly, the algorithm successively inflates balls with radius r around each point of the manifold. If k points have a pairwise distance smaller than r (that is, for all pairs of points both points are contained in the ball of the other point), they form a $(k-1)$D simplex. Thus, the neural manifold gives rise to a simplicial complex (a collection of simplices of potentially different dimension) the topological features of which represent the topology of the neural manifold and can be extracted computationally. As r is increased, many short-lived features appear by chance. If the manifold has complex topological structures, they should continuously appear as the radius of the balls grows for a large range of r. We computed the persistence barcode for the first three homology groups H_0, H_1 and H_2. Homology groups are topological invariants that capture topological features of a given dimension of the neural manifold. The long-lasting bars in the n-th persistence barcode correspond to the number of independent generators β_n of the respective homology group H_n. For low dimensionalities, they can be interpreted intuitively: β_0 is the number of connected components, β_1 the number of 1D holes, β_2 the number of enclosed 2D voids. Throughout all plots of this paper we display the top 1% longest-lasting barcodes for each homology group.

Dimensionality

We used two different approaches to study the dimensionality of the neural data.

First, we compute the time-varying participation ratio (PR, Equation 1) from the covariance matrix. We take a 30 s sliding window with a 1 s offset over the MUAe data and compute the PR for each window separately. Higher activity leads to higher variance; thus, we normalised the data within each window via z-scoring to minimise this effect. The PR does not require setting an arbitrary threshold. From the time-varying PR we measured the correlation between the log odds and the PR, and the PR distribution in each manifold.

Second, we computed the eigenvalue distribution of the neural data for the entire session as well as within each manifold. Once again we normalised the data after sampling each manifold. The distribution appeared to follow a power law, in agreement with previous studies[10]. We used a linear regression in log-log space to fit a power law to our data, where the slope of the linear fit in the log-log plot corresponds to the exponent α of the power law.

Coherence and Granger causality

To estimate the communication between cortical areas we rely on the coherence and Granger causality.

Coherence is the quantification of linear correlations in the frequency domain. Such that

$$C_{xy}(f) = \frac{|S_{xy}(f)|^2}{S_{xx}(f)S_{yy}(f)} \tag{2}$$

[2]https://pypi.org/project/ripser/

where C_{xy} is the frequency (f) dependent coherence between two signals x and y, $S_{xy}(f)$ is the cross-spectral density, and $S_{xx}(f)$ and $S_{yy}(f)$ are the auto-spectral densities.

In order to assess the directionality of frequency dependent interactions between the areas we applied spectral Granger causality analysis to the LFP recordings[71]. We first computed the cross-spectral matrix $S(f)$ with the multitaper method. To this end, we subdivided the chosen signal pairs into 10 s long segments. These were processed individually with 3 Slepian tapers and averaged in the end. This yielded the cross-spectrum. The segments had an overlap of 50%. Next, we decomposed the cross-spectrum into the covariance matrix Σ and the transfer function $H(f)$ with the Wilson spectral matrix factorisation[72], obtaining the matrix equation

$$S = H(f)\Sigma H^{\dagger}(f). \tag{3}$$

With these factors, one is able to obtain a version of directional functional connectivity between the first and second signals via

$$\mathrm{GC}_{x\to y}(f) = \frac{S_{xx}(f)}{\tilde{H}_{xx}(f)\Sigma_{xx}\tilde{H}^{\dagger}_{xx}(f)} \tag{4}$$

where $\tilde{H}_{xx}(f) = H_{xx}(f) + \Sigma_{xy}/\Sigma_{xx}H_{xy}(f)$ and mutatis mutandis for the influence of the second onto the first signal. The analysis was performed for all pairs of channels between the areas that exhibited a peak in the coherence in the β band 12 Hz $< f <$ 30 Hz.

We quantify the beta-band Granger causality strength as

$$\mathcal{B} = \sum_{f=12\ \mathrm{Hz}}^{f=30\ \mathrm{Hz}} \mathrm{GC}_{x\leftarrow y}(f) - \mathrm{GC}_{x\to y}(f). \tag{5}$$

We also analysed the time-varying spectral Granger causality. For this aim we used 10 s windows and moved them across the data with 1 s steps, for a final time resolution of 1 Hz. We calculated the spectral Granger causality for each window separately. The initial and final 5 s were discarded to avoid disruptions at the boundaries. So the time-varying causality spectrogram is

$$\mathrm{GC}_{x\to y}(t,f) = \mathrm{GC}_{x\to y}(f)\Big|_{[t_0,t_1]}, \ \dots\ , \mathrm{GC}_{x\to y}(f)\Big|_{[t_{n-1},t_n]} \tag{6}$$

and mutatis mutandis for the $y \to x$ direction.

Finally, we also define the time-varying Granger causality difference

$$\Delta\mathrm{GC}(t,f) = \mathrm{GC}_{x\leftarrow y}(t,f) - \mathrm{GC}_{x\to y}(t,f), \tag{7}$$

which summed over the beta band we call

$$\mathcal{B}(t) = \sum_{f=12\ \mathrm{Hz}}^{f=30\ \mathrm{Hz}} \Delta\mathrm{GC}(t,f). \tag{8}$$

Note that

$$\mathcal{B} \neq \frac{\sum_{i}^{N} \mathcal{B}(t_i)}{N} \tag{9}$$

due to the nonlinearities in the Granger causality calculation.

Both the coherence and spectral Granger causality were implemented in the Electrophysiology Analysis Toolkit (Elephant)[73].

Peak detection

Table 2: Peak detection algorithm parameters.

CWT peak detection parameters		
widths	100 – 500 Hz	Width range for CWT matrix.
width_step	0.1 Hz	Step between widths.
wavelet	Ricker	Wavelet used for convolution.
max_distances	widths / 4	Criterion to consider ridge lines connected.
gap_thresh	10 Hz	Ridge lines farther apart will not be connected.
min_length	225	Minimum length of ridge lines.
min_snr	1	Minimum snr of ridge lines.
noise_perc	10	Percentile of ridge line considered noise for snr calculation.

To detect coherence peaks we used a standard peak detection algorithm for 1D arrays using wavelet transforms. We computed the continuous wavelet transform (cwt) for wavelets with widths from 10 Hz to 100 Hz (at 0.1 Hz steps), using a Ricker wavelet—i.e., a Mexican hat. Next, we searched for ridge lines in the cwt—peaks across different wavelet lengths—following standard criteria[74]. Finally, the ridge lines were filtered based on their total length, gaps, and signal-to-noise ratio (snr). The resulting ridge lines (if any) were considered as peaks in the coherence.

The detected peaks tended to be broad, since our parameter choice intentionally rejected narrow peaks. We chose this configuration in favour of robustness and to minimise false positives. Nevertheless, peaks were detected for a majority of electrode pairs.

Spiking neural network simulations

Table 3: General model description

Model Summary	
Populations	two populations, one excitatory, one inhibitory
Connectivity	random connectivity
Neuron model	leaky integrate-and-fire model
Synapse model	exponential postsynaptic current
Input	independent spike trains from inhomogeneous Poisson processes with given rate $r(t)$
Neuron and synapse model	
Subthreshold dynamics	$\frac{dV}{dt} = -\frac{V}{\tau_m} + \frac{I_{syn}(t)}{C_m}$, $I_{syn}(t) = Je^{-(t-t^*-d)/\tau_{syn}} H(t-t^*-d)$
Spiking	If $V(t-) < \theta$ and $V(t+) \geq \theta$, 1. Set $t^* = t$ and $V(t) = V_0$, and 2. Emit spike with time stamp t^*.
Connectivity	
Type	pairwise Bernoulli, i.e., for each pair of neurons generate a synapse with probability p
Weights	fixed source- and target-population-specific weights
Delays	log-normally distributed delays for excitatory and inhibitory neurons
Input	
Background	$r(t) = \max(0, \nu^{\text{bg}}_{\text{base}} + \nu^{\text{bg}}_{\text{amp}} \cdot \sin(2\pi f^{\text{bg}} \cdot t))$
Top-down modulation	$r(t) = \max(0, \nu^{\text{td}}_{\text{base}} + \nu^{\text{td}}_{\text{amp}} \cdot \sin(2\pi f^{\text{td}} \cdot t))$

To investigate the hypothesis that top-down signals in the β-band induce a change in the population dynamics and dimensionality, we conducted a spiking neural network simulation. The network consists of $10,000$ excitatory and $2,000$ inhibitory leaky integrate-and-fire (LIF) neurons with exponential post-synaptic currents. Pairs of neurons are randomly connected with a connection probability of $p = 0.1$. The spike transmission delay is randomly sampled following a log-normal distribution. Generally speaking, the simulation experiments consist of two parts corresponding to the two states observed in the neuronal activity. In the first state (background state), the input consists of spike trains sampled from an inhomogeneous Poisson process with a baseline rate of ν_{bg} Hz that is modulated with a 1 Hz sinusoidal oscillation. In the second state, the network additionally receives input spike trains from inhomogeneous Poisson processes with rates oscillating at 20 Hz. The first state represents the eyes-closed, the second the eyes-open condition. Both input regimes provide independent input to each neuron, based on the same rate profiles. During the simulation, we recorded the spiking activity of $1,000$ excitatory and 200 inhibitory neurons. We provided the top-down modulation to a subset of the neurons in the network. We targeted 50% of both the excitatory and inhibitory population. During the simulation, we distinguish two manifolds corresponding to the eyes-open and eyes-closed periods during the recordings. See Table 3, Table 4 for a full description of the network and the experiments. For the simulations we used NEST (version 3.3) [75].

Table 4: Simulation parameters.

Population Parameters		
N_{ex}	10,000	number of excitatory neurons
N_{in}	2,000	number of inhibitory neurons
Connectivity Parameters		
p	0.1	connection probability
Neuron parameters		
τ_{m}	20 ms	membrane time constant
τ_{r}	2 ms	absolute refractory period
τ_{syn}	2 ms	postsynaptic current time constant
C_{m}	250 pF	membrane capacity
V_{m}	0 mV	resting potential
E_{L}	0 mV	membrane capacity
V_{reset}	0 mV	reset membrane potential
V_{th}	20 mV	threshold
Stimulus parameters: Background		
$\nu_{\mathrm{base}}^{\mathrm{bg}}$	8682 spikes/s	baseline rate
$\nu_{\mathrm{amp}}^{\mathrm{bg}}$	2170 spikes/s	amplitude
f	10 Hz	sinusoidal oscillation frequency
Stimulus parameters: Top-down signal		
$\nu_{\mathrm{base}}^{\mathrm{td}}$	0 spikes/s	base line rate
$\nu_{\mathrm{amp}}^{\mathrm{td}}$	723 spikes/s	amplitude
f	20 Hz	sinusoidal oscillation frequency
p^{td}	0.5	fraction of neurons targeted by top-down modulation in setup 1 and 3
p^{td}	1	fraction of neurons targeted by top-down modulation in setup 2
Synapse parameters		
J_{EE}	6.4 pA	synaptic efficacy excitatory to excitatory
J_{IE}	9.5 pA	synaptic efficacy excitatory to inhibitory
g	4	relative inhibitory synaptic efficacy
J_{EI}	$-g * J_{EE}$	synaptic efficacy inhibitory to excitatory
J_{II}	$-g * J_{EE}$	synaptic efficacy inhibitory to inhibitory
Delay parameters		
μ_{ex}	1.5 ms	mean of lognormal distribution for excitatory connections
μ_{in}	0.75 ms	mean of lognormal distribution for inhibitory connections
$\sigma_{\mathrm{ex,in}}$	0.5 ms	standard deviation of lognormal distribution for all connections

Author contributions

AMG, JI, and AK conceptualised the study. AMG curated and processed the data from macaques L and A. AKJ curated and processed the data from macaque Y, with minor inputs from AMG. AMG and AK performed the dimensionality analysis, statistical testing, and Granger causality analysis. AK performed the spiking neuron network simulations and AMG analyzed the results. AMG created all figures, with feedback from all other authors. AMG wrote the initial manuscript; all other authors provided feedback. JI, SG, and SvA supervised the project. SG and SvA procured the funding.

Acknowledgements

We thank David Dahmen for his support with the dimensionality analysis and modelling. We thank Simon Essink for useful feedback.

This project received funding from the DFG Priority Program (SPP 2041 "Computational Connectomics") [S.J.van Albada: AL 2041/1-1]; the EU's Horizon 2020 Framework Grant Agreement No. 945539 (Human Brain Project SGA3); the DFG (RTG 2416 "MultiSenses-MultiScales"); and the CNRS Multidisciplinary Exploratory Projects initiative.

Conflict of Interest

The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

References

[1] Peiran Gao, Eric Trautmann, Byron Yu, Gopal Santhanam, Stephen Ryu, Krishna Shenoy, and Surya Ganguli. A theory of multineuronal dimensionality, dynamics and measurement. *bioRxiv*, 2017. doi: 10.1101/214262. URL https://www.biorxiv.org/content/10.1101/214262v2.

[2] Juan A. Gallego, Matthew G. Perich, Lee E. Miller, and Sara A. Solla. Neural Manifolds for the Control of Movement. *Neuron*, 94(5):978–984, 2017. ISSN 0896-6273. doi: 10.1016/j.neuron.2017.05.025. URL https://www.cell.com/neuron/abstract/S0896-6273(17)30463-4.

[3] Juan A. Gallego, Matthew G. Perich, Stephanie N. Naufel, Christian Ethier, Sara A. Solla, and Lee E. Miller. Cortical population activity within a preserved neural manifold underlies multiple motor behaviors. *Nature Communications*, 9(1):1–13, 2018. ISSN 2041-1723. doi: 10.1038/s41467-018-06560-z. URL https://www.nature.com/articles/s41467-018-06560-z.

[4] Barbara Feulner and Claudia Clopath. Neural manifold under plasticity in a goal driven learning behaviour. *PLOS Computational Biology*, 17(2):e1008621, February 2021. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1008621. URL https://dx.plos.org/10.1371/journal.pcbi.1008621.

[5] Rishidev Chaudhuri, Berk Gerçek, Biraj Pandey, Adrien Peyrache, and Ila Fiete. The intrinsic attractor manifold and population dynamics of a canonical cognitive circuit across waking and sleep. *Nature Neuroscience*, 22(9): 1512–1520, 2019. ISSN 1546-1726. doi: 10.1038/s41593-019-0460-x. URL https://www.nature.com/articles/s41593-019-0460-x.

[6] Valerio Mante, David Sussillo, Krishna V. Shenoy, and William T. Newsome. Context-dependent computation by recurrent dynamics in prefrontal cortex. *Nature*, 503(7474):78–84, November 2013. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature12742. URL http://www.nature.com/articles/nature12742.

[7] Mark M. Churchland, John P. Cunningham, Matthew T. Kaufman, Justin D. Foster, Paul Nuyujukian, Stephen I. Ryu, and Krishna V. Shenoy. Neural population dynamics during reaching. *Nature*, 487(7405):51–56, July 2012. ISSN 0028-0836, 1476-4687. doi: 10.1038/nature11129. URL http://www.nature.com/articles/nature11129.

[8] Anqi Wu, Stan Pashkovski, Sandeep R Datta, and Jonathan W Pillow. Learning a latent manifold of odor representations from neural responses in piriform cortex. In S. Bengio, H. Wallach, H. Larochelle, K. Grauman, N. Cesa-Bianchi, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 31, 2018. URL https://proceedings.neurips.cc/paper/2018/file/17b3c7061788dbe82de5abe9f6fe22b3-Paper.pdf.

[9] Richard J. Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A. Baas, Benjamin A. Dunn, May-Britt Moser, and Edvard I. Moser. Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128, 2022. ISSN 0028-0836, 1476-4687. doi: 10.1038/s41586-021-04268-7. URL https://www.nature.com/articles/s41586-021-04268-7.

[10] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Matteo Carandini, and Kenneth D. Harris. High-dimensional geometry of population responses in visual cortex. *Nature*, 571(7765):361–365, 2019. ISSN 1476-4687. doi: 10.1038/s41586-019-1346-5. URL https://www.nature.com/articles/s41586-019-1346-5.

[11] Emmanouil Froudarakis, Uri Cohen, Maria Diamantaki, Edgar Y. Walker, Jacob Reimer, Philipp Berens, Haim Sompolinsky, and Andreas S. Tolias. Object manifold geometry across the mouse cortical visual hierarchy. *bioRxiv*, 2020. doi: 10.1101/2020.08.20.258798. URL http://biorxiv.org/lookup/doi/10.1101/2020.08.20.258798.

[12] Gurjeet Singh, Facundo Memoli, Tigran Ishkhanov, Guillermo Sapiro, Gunnar Carlsson, and Dario L. Ringach. Topological analysis of population activity in visual cortex. *Journal of Vision*, 8(8):11–11, 2008. ISSN 1534-7362. doi: 10.1167/8.8.11. URL https://jov.arvojournals.org/article.aspx?articleid=2193262.

[13] Luca Baroni, Mohammad Bashiri, Konstantin Friedrich Willeke, Ján Antolík, and Fabian H. Sinz. Learning invariance manifolds of visual sensory neurons. In *NeurIPS 2022 Workshop on Symmetry and Geometry in Neural Representations*, 2022. URL https://openreview.net/forum?id=2dQyENiU330.

[14] Robert Ghrist. Barcodes: The persistent topology of data. *Bulletin of the American Mathematical Society*, 45(1): 61–75, 2008. ISSN 0273-0979, 1088-9485. doi: 10.1090/S0273-0979-07-01191-3. URL https://www.ams.org/bull/2008-45-01/S0273-0979-07-01191-3/.

[15] John P Cunningham and Byron M Yu. Dimensionality reduction for large-scale neural recordings. *Nature Neuroscience*, 17(11):1500–1509, November 2014. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn.3776. URL http://www.nature.com/articles/nn.3776.

[16] Peiran Gao and Surya Ganguli. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current Opinion in Neurobiology*, 32:148–155, 2015. ISSN 09594388. doi: 10.1016/j.conb.2015.04.003. URL https://linkinghub.elsevier.com/retrieve/pii/S0959438815000768.

[17] Rich Pang, Benjamin J. Lansdell, and Adrienne L. Fairhall. Dimensionality reduction in neuroscience. *Current Biology*, 26(14):R656–R660, July 2016. ISSN 09609822. doi: 10.1016/j.cub.2016.05.029. URL https://linkinghub.elsevier.com/retrieve/pii/S0960982216304870.

[18] Ryan C. Williamson, Benjamin R. Cowley, Ashok Litwin-Kumar, Brent Doiron, Adam Kohn, Matthew A. Smith, and Byron M. Yu. Scaling Properties of Dimensionality Reduction for Neural Populations and Network Models. *PLOS Computational Biology*, 12(12):e1005141, December 2016. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1005141. URL https://dx.plos.org/10.1371/journal.pcbi.1005141.

[19] Luca Mazzucato, Alfredo Fontanini, and Giancarlo La Camera. Stimuli Reduce the Dimensionality of Cortical Activity. *Frontiers in Systems Neuroscience*, 10, February 2016. ISSN 1662-5137. doi: 10.3389/fnsys.2016.00011. URL http://journal.frontiersin.org/Article/10.3389/fnsys.2016.00011/abstract.

[20] Lilach Avitan and Carsen Stringer. Not so spontaneous: Multi-dimensional representations of behaviors and context in sensory areas. *Neuron*, page S0896627322005888, July 2022. ISSN 08966273. doi: 10.1016/j.neuron.2022.06.019. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627322005888.

[21] Esther Marx, Thomas Stephan, Annina Nolte, Angela Deutschländer, Klaus C Seelos, Marianne Dieterich, and Thomas Brandt. Eye closure in darkness animates sensory systems. *NeuroImage*, 19(3):924–934, 2003. ISSN 10538119. doi: 10.1016/S1053-8119(03)00150-2. URL https://linkinghub.elsevier.com/retrieve/pii/S1053811903001502.

[22] Esther Marx, Angela Deutschländer, Thomas Stephan, Marianne Dieterich, Martin Wiesmann, and Thomas Brandt. Eyes open and eyes closed as rest conditions: impact on brain activation patterns. *NeuroImage*, 21(4):1818–1824, 2004. ISSN 10538119. doi: 10.1016/j.neuroimage.2003.12.026. URL https://linkinghub.elsevier.com/retrieve/pii/S1053811903007936.

[23] Yu. A. Boytsova and S. G. Danko. EEG differences between resting states with eyes open and closed in darkness. *Human Physiology*, 36(3):367–369, 2010. ISSN 0362-1197, 1608-3164. doi: 10.1134/S0362119710030199. URL http://link.springer.com/10.1134/S0362119710030199.

[24] Mark McAvoy, Linda Larson-Prior, Marek Ludwikow, Dongyang Zhang, Abraham Z. Snyder, Debra L. Gusnard, Marcus E. Raichle, and Giovanni d'Avossa. Dissociated mean and functional connectivity BOLD signals in visual cortex during eyes closed and fixation. *Journal of Neurophysiology*, 108(9):2363–2372, 2012. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00900.2011. URL https://www.physiology.org/doi/10.1152/jn.00900.2011.

[25] Stefan Brodoehl, Carsten M. Klingner, and Otto W. Witte. Eye closure enhances dark night perceptions. *Scientific Reports*, 5(1):10515, 2015. ISSN 2045-2322. doi: 10.1038/srep10515. URL http://www.nature.com/articles/srep10515.

[26] Hans Berger. Über das Elektrenkephalogramm des Menschen. *Archiv für Psychiatrie und Nervenkrankheiten*, 87(1): 527–570, December 1929. ISSN 0003-9373, 1433-8491. doi: 10.1007/BF01797193. URL http://link.springer.com/10.1007/BF01797193.

[27] Wolfgang Klimesch, Paul Sauseng, and Simon Hanslmayr. EEG alpha oscillations: The inhibition–timing hypothesis. *Brain Research Reviews*, 53(1):63–88, 2007. ISSN 01650173. doi: 10.1016/j.brainresrev.2006.06.003.

[28] David T.J. Liley and Suresh D. Muthukumaraswamy. Evidence that alpha blocking is due to increases in system-level oscillatory damping not neuronal population desynchronisation. *NeuroImage*, 208:116–408, March 2020. ISSN 10538119. doi: 10.1016/j.neuroimage.2019.116408.

[29] Xing Chen, Feng Wang, Eduardo Fernandez, and Pieter R. Roelfsema. Shape perception via a high-channel-count neuroprosthesis in monkey visual cortex. *Science*, 370(6521):1191–1196, 2020. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.abd7435. URL https://www.science.org/doi/10.1126/science.abd7435.

[30] Timo van Kerkoerle, Matthew W. Self, Bruno Dagnino, Marie-Alice Gariel-Mathis, Jasper Poort, Chris van der Togt, and Pieter R. Roelfsema. Alpha and gamma oscillations characterize feedback and feedforward processing in monkey visual cortex. *Proceedings of the National Academy of Sciences*, 111(40):14332–14341, 2014. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1402773111. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1402773111.

[31] André Moraes Bastos, Julien Vezoli, Conrado Arturo Bosman, Jan-Mathijs Schoffelen, Robert Oostenveld, Jarrod Robert Dowdall, Peter De Weerd, Henry Kennedy, and Pascal Fries. Visual Areas Exert Feedforward and Feedback Influences through Distinct Frequency Channels. *Neuron*, 85(2):390–401, 2015. ISSN 08966273. doi: 10.1016/j.neuron.2014.12.018. URL https://linkinghub.elsevier.com/retrieve/pii/S089662731401099X.

[32] P. Christiaan Klink, Bruno Dagnino, Marie-Alice Gariel-Mathis, and Pieter R. Roelfsema. Distinct Feedforward and Feedback Effects of Microstimulation in Visual Cortex Reveal Neural Mechanisms of Texture Segregation. *Neuron*, 95 (1):209–220.e3, 2017. ISSN 08966273. doi: 10.1016/j.neuron.2017.05.033. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627317304713.

[33] Jasper Poort, Florian Raudies, Aurel Wannig, Victor A.F. Lamme, Heiko Neumann, and Pieter R. Roelfsema. The Role of Attention in Figure-Ground Segregation in Areas V1 and V4 of the Visual Cortex. *Neuron*, 75(1):143–156, July 2012. ISSN 08966273. doi: 10.1016/j.neuron.2012.04.032. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627312004369.

[34] Matthew W. Self, Timo van Kerkoerle, Hans Supèr, and Pieter R. Roelfsema. Distinct Roles of the Cortical Layers of Area V1 in Figure-Ground Segregation. *Current Biology*, 23(21):2121–2129, 2013. ISSN 09609822. doi: 10.1016/j.cub.2013.09.013. URL https://linkinghub.elsevier.com/retrieve/pii/S0960982213011299.

[35] Timo van Kerkoerle, Matthew W. Self, and Pieter R. Roelfsema. Layer-specificity in the effects of attention and working memory on activity in primary visual cortex. *Nature Communications*, 8(1):13804, 2017. ISSN 2041-1723. doi: 10.1038/ncomms13804. URL http://www.nature.com/articles/ncomms13804.

[36] Hualou Liang, Xiajing Gong, Minggui Chen, Yin Yan, Wu Li, and Charles D. Gilbert. Interactions between feedback and lateral connections in the primary visual cortex. *Proceedings of the National Academy of Sciences*, 114(32): 8637–8642, 2017. ISSN 0027-8424, 1091-6490. doi: 10.1073/pnas.1706183114. URL http://www.pnas.org/lookup/doi/10.1073/pnas.1706183114.

[37] Laura B Naumann, Joram Keijser, and Henning Sprekeler. Invariant neural subspaces maintained by feedback modulation. *eLife*, 11:e76096, 2022. ISSN 2050-084X. doi: 10.7554/eLife.76096. URL https://elifesciences.org/articles/76096.

[38] David Dahmen, Stefano Recanatesi, Xiaoxuan Jia, Gabriel K. Ocker, Luke Campagnola, Tim Jarsky, Stephanie Seeman, Moritz Helias, and Eric Shea-Brown. Strong and localized recurrence controls dimensionality of neural activity across brain areas. *bioRxiv*, 2020. doi: 10.1101/2020.11.02.365072. URL http://biorxiv.org/lookup/doi/10.1101/2020.11.02.365072.

[39] Ad Aertsen and Hubert Preissl. Dynamics of activity and connectivity in physiological neuronal networks. *Nonlinear Dynamics and Neuronal Networks*, pages 281–301, 1991.

[40] Hans Supèr and Pieter R. Roelfsema. Chronic multiunit recordings in behaving animals: advantages and limitations. *Progress in Brain Research*, 147:263–282, 2005. doi: 10.1016/S0079-6123(04)47020-4.

[41] Xing Chen, Aitor Morales-Gregorio, Julia Sprenger, Alexander Kleinjohann, Shashwat Sridhar, Sacha J. van Albada, Sonja Grün, and Pieter R. Roelfsema. 1024-channel electrophysiological recordings in macaque V1 and V4 during resting state. *Scientific Data*, 9(1):77, 2022. ISSN 2052-4463. doi: 10.1038/s41597-022-01180-1. URL https://www.nature.com/articles/s41597-022-01180-1.

[42] Marcel Jan de Haan, Thomas Brochier, Sonja Grün, Alexa Riehle, and Frédéric V. Barthélemy. Real-time visuomotor behavior and electrophysiology recording setup for use with humans and monkeys. *Journal of Neurophysiology*, 120 (2):539–552, August 2018. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.00262.2017. URL https://www.physiology.org/doi/10.1152/jn.00262.2017.

[43] Mingli Wang, Yujie Hou, Loïc Magrou, Joonas A. Autio, Pierre Misery, Tim Coalson, Erin Reid, Yuanfang Xu, Camille Lamy, Arnauld Falchier, Qi Zhang, Mu-Ming Poo, Colette Dehay, Matthew F. Glasser, Takuya Hayashi, Kenneth Knoblauch, David Van Essen, Zhiming Shen, and Henry Kennedy. Retinotopic organization of feedback projections in primate early visual cortex: implications for active vision. *bioRxiv*, 2022. doi: 10.1101/2022.04.27.489651. URL http://biorxiv.org/lookup/doi/10.1101/2022.04.27.489651.

[44] Julien Vezoli, Martin Vinck, Conrado Arturo Bosman, André Moraes Bastos, Christopher Murphy Lewis, Henry Kennedy, and Pascal Fries. Brain rhythms define distinct interaction networks with differential dependence on anatomy. *Neuron*, 109(23):3862–3878.e5, December 2021. ISSN 08966273. doi: 10.1016/j.neuron.2021.09.052. URL https://linkinghub.elsevier.com/retrieve/pii/S089662732100725X.

[45] Carsen Stringer, Marius Pachitariu, Nicholas Steinmetz, Charu Bai Reddy, Matteo Carandini, and Kenneth D. Harris. Spontaneous behaviors drive multidimensional, brainwide activity. *Science*, 364(6437):eaav7893, 2019. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.aav7893.

[46] Bharath C. Talluri, Incheol Kang, Adam Lazere, Katrina R. Quinn, Nicholas Kaliss, Jacob L. Yates, Daniel A. Butts, and Hendrikje Nienborg. Activity in primate visual cortex is minimally driven by spontaneous movements. *bioRxiv*, 2022. doi: 10.1101/2022.09.08.507006.

[47] Paulina Anna Dabrowska, Nicole Voges, Michael von Papen, Junji Ito, David Dahmen, Alexa Riehle, Thomas Brochier, and Sonja Grün. On the Complexity of Resting State Spiking Activity in Monkey Motor Cortex. *Cerebral Cortex Communications*, 2(3):tgab033, 2021. ISSN 2632-7376. doi: 10.1093/texcom/tgab033. URL https://academic.oup.com/cercorcomms/article/doi/10.1093/texcom/tgab033/6277047.

[48] Diego Mendoza-Halliday, Alex J. Major, Noah Lee, Maxwell Lichtenfeld, Brock Carlson, Blake Mitchell, Patrick D. Meng, Yihan (Sophy) Xiong, Jacob A. Westerberg, Alexander Maier, Robert Desimone, Earl K. Miller, and André M. Bastos. A ubiquitous spectrolaminar motif of local field potential power across the primate cortex. *bioRxiv*, October 2022. doi: 10.1101/2022.09.30.510398. URL http://biorxiv.org/lookup/doi/10.1101/2022.09.30.510398.

[49] João D. Semedo, Anna I. Jasper, Amin Zandvakili, Aravind Krishna, Amir Aschner, Christian K. Machens, Adam Kohn, and Byron M. Yu. Feedforward and feedback interactions between visual cortical areas use different population activity patterns. *Nature Communications*, 13(1):1099, December 2022. ISSN 2041-1723. doi: 10.1038/s41467-022-28552-w. URL https://www.nature.com/articles/s41467-022-28552-w.

[50] Nikola T. Markov, Julien Vezoli, Pascal Chameau, Arnaud Falchier, René Quilodran, Cyril Huissoud, Camille Lamy, Pierre Misery, Pascale Giroud, Shimon Ullman, Pascal Barone, Colette Dehay, Kenneth Knoblauch, and Henry Kennedy. Anatomy of hierarchy: Feedforward and feedback pathways in macaque visual cortex. *Journal of Comparative Neurology*, 522(1):225–259, 2014. ISSN 00219967. doi: 10.1002/cne.23458.

[51] Stephen G. Lomber, Bertram R. Payne, and James A. Horel. The cryoloop: an adaptable reversible cooling deactivation method for behavioral or electrophysiological assessment of neural function. *Journal of Neuroscience Methods*, 86 (2):179–194, 1999. ISSN 01650270. doi: 10.1016/S0165-0270(98)00165-4. URL https://linkinghub.elsevier.com/retrieve/pii/S0165027098001654.

[52] J. J. Nassi, S. G. Lomber, and R. T. Born. Corticocortical Feedback Contributes to Surround Suppression in V1 of the Alert Primate. *Journal of Neuroscience*, 33(19):8504–8517, May 2013. ISSN 0270-6474, 1529-2401. doi: 10.1523/JNEUROSCI.5124-12.2013.

[53] E. C. Dias, M. Kiesau, and M. A. Segraves. Acute activation and inactivation of macaque frontal eye field with GABA-related drugs. *Journal of Neurophysiology*, 74(6):2744–2748, December 1995. ISSN 0022-3077, 1522-1598. doi: 10.1152/jn.1995.74.6.2744. URL https://www.physiology.org/doi/10.1152/jn.1995.74.6.2744.

[54] Thomas Brochier, Marie-Josée Boudreau, Michel Paré, and A. M. Smith. The effects of muscimol inactivation of small regions of motor and somatosensory cortex on independent finger movements and force control in the precision grip. *Experimental Brain Research*, 128(1-2):31–40, September 1999. ISSN 00144819. doi: 10.1007/s002210050814. URL http://link.springer.com/10.1007/s002210050814.

[55] Longtang L Chen, Laurent Goffart, and David L Sparks. A simple method for constructing microinjectrodes for reversible inactivation in behaving monkeys. *Journal of Neuroscience Methods*, 107(1-2):81–85, May 2001. ISSN 01650270. doi: 10.1016/S0165-0270(01)00354-5.

[56] Samantha R. Debes and Valentin Dragoi. Suppressing feedback signals to visual cortex abolishes attentional modulation. *Science*, 379(6631):468–473, 2023. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.ade1855.

[57] Nicolas Brunel. Dynamics of Sparsely Connected Networks of Excitatory and Inhibitory Spiking Neurons. *Journal of Computational Neuroscience*, 8(3):183–208, 2000. ISSN 09295313. doi: 10.1023/A:1008925309027. URL http://link.springer.com/10.1023/A:1008925309027.

[58] Brent Doiron, Ashok Litwin-Kumar, Robert Rosenbaum, Gabriel K Ocker, and Krešimir Josić. The mechanics of state-dependent neural correlations. *Nature Neuroscience*, 19(3):383–393, 2016. ISSN 1097-6256, 1546-1726. doi: 10.1038/nn.4242. URL http://www.nature.com/articles/nn.4242.

[59] Tom Tetzlaff, Moritz Helias, Gaute T. Einevoll, and Markus Diesmann. Decorrelation of Neural-Network Activity by Inhibitory Feedback. *PLoS Computational Biology*, 8(8):e1002596, 2012. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1002596. URL https://dx.plos.org/10.1371/journal.pcbi.1002596.

[60] Michael T. Schaub, Yazan N. Billeh, Costas A. Anastassiou, Christof Koch, and Mauricio Barahona. Emergence of Slow-Switching Assemblies in Structured Neuronal Networks. *PLOS Computational Biology*, 11(7):e1004196, 2015. ISSN 1553-7358. doi: 10.1371/journal.pcbi.1004196. URL https://dx.plos.org/10.1371/journal.pcbi.1004196.

[61] Vahid Rostami, Thomas Rost, Alexa Riehle, Sacha J. Van Albada, and Martin P. Nawrot. Excitatory and inhibitory motor cortical clusters account for balance, variability, and task performance. preprint, Neuroscience, 2020. URL http://biorxiv.org/lookup/doi/10.1101/2020.02.27.968339.

[62] Anouk M. Van Loon, Johannes J. Fahrenfort, Bauke Van Der Velde, Philipp B. Lirk, Nienke C. C. Vulink, Markus W. Hollmann, H. Steven Scholte, and Victor A. F. Lamme. NMDA Receptor Antagonist Ketamine Distorts Object Recognition by Reducing Feedback to Early Visual Cortex. *Cerebral Cortex*, 26(5):1986–1996, May 2016. ISSN 1047-3211, 1460-2199. doi: 10.1093/cercor/bhv018. URL https://academic.oup.com/cercor/article-lookup/doi/10.1093/cercor/bhv018.

[63] Jose L. Herrero, Marc A. Gieselmann, Mehdi Sanayei, and Alexander Thiele. Attention-Induced Variance and Noise Correlation Reduction in Macaque V1 Is Mediated by NMDA Receptors. *Neuron*, 78(4):729–739, May 2013. ISSN 08966273. doi: 10.1016/j.neuron.2013.03.029. URL https://linkinghub.elsevier.com/retrieve/pii/S0896627313002766.

[64] John D. Porter, Barton L. Guthrie, and David L. Sparks. Innervation of monkey extraocular muscles: Localization of sensory and motor neurons by retrograde transport of horseradish peroxidase. *The Journal of Comparative Neurology*, 218(2):208–219, August 1983. ISSN 0021-9967, 1096-9861. doi: 10.1002/cne.902180208. URL https://onlinelibrary.wiley.com/doi/10.1002/cne.902180208.

[65] Daniela Balslev, Neil B. Albert, and Chris Miall. Eye muscle proprioception is represented bilaterally in the sensorimotor cortex. *Human Brain Mapping*, 32(4):624–631, April 2011. ISSN 1065-9471, 1097-0193. doi: 10.1002/hbm.21050. URL https://onlinelibrary.wiley.com/doi/10.1002/hbm.21050.

[66] N. T. Markov, M. M. Ercsey-Ravasz, A. R. Ribeiro Gomes, C. Lamy, L. Magrou, J. Vezoli, P. Misery, A. Falchier, R. Quilodran, M. A. Gariel, J. Sallet, R. Gamanut, C. Huissoud, S. Clavagnier, P. Giroud, D. Sappey-Marinier, P. Barone, C. Dehay, Z. Toroczkai, K. Knoblauch, D. C. Van Essen, and H. Kennedy. A Weighted and Directed Interareal Connectivity Matrix for Macaque Cerebral Cortex. *Cerebral Cortex*, 24(1):17–36, January 2014. ISSN 1460-2199, 1047-3211. doi: 10.1093/cercor/bhs270.

[67] Georgia G. Gregoriou, Stephen J. Gotts, Huihui Zhou, and Robert Desimone. High-Frequency, Long-Range Coupling Between Prefrontal and Visual Cortex During Attention. *Science*, 324(5931):1207–1210, 2009. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.1171402.

[68] Simon Hanslmayr, Wolfgang Klimesch, Paul Sauseng, Walter Gruber, Michael Doppelmayr, Roman Freunberger, and Thomas Pecherstorfer. Visual discrimination performance is related to decreased alpha amplitude but increased phase locking. *Neuroscience Letters*, 375(1):64–68, 2005. ISSN 03043940. doi: 10.1016/j.neulet.2004.10.092. URL https://linkinghub.elsevier.com/retrieve/pii/S030439400401359X.

[69] Pierre Yger, Giulia LB Spampinato, Elric Esposito, Baptiste Lefebvre, Stéphane Deny, Christophe Gardella, Marcel Stimberg, Florian Jetter, Guenther Zeck, Serge Picaud, Jens Duebel, and Olivier Marre. A spike sorting toolbox for up to thousands of electrodes validated with ground truth recordings in vitro and in vivo. *eLife*, 7:e34518, 2018. ISSN 2050-084X. doi: 10.7554/eLife.34518. URL https://elifesciences.org/articles/34518.

[70] Joshua B. Tenenbaum, Vin de Silva, and John C. Langford. A Global Geometric Framework for Nonlinear Dimensionality Reduction. *Science*, 290(5500):2319–2323, December 2000. ISSN 0036-8075, 1095-9203. doi: 10.1126/science.290.5500.2319. URL https://www.science.org/doi/10.1126/science.290.5500.2319.

[71] Mukeshwar Dhamala, Govindan Rangarajan, and Mingzhou Ding. Estimating Granger Causality from Fourier and Wavelet Transforms of Time Series Data. *Physical Review Letters*, 100(1):018701, 2008. ISSN 0031-9007, 1079-7114. doi: 10.1103/PhysRevLett.100.018701. URL https://link.aps.org/doi/10.1103/PhysRevLett.100.018701.

[72] G. Tunnicliffe Wilson. The Factorization of Matricial Spectral Densities. *SIAM Journal on Applied Mathematics*, 23 (4):420–426, 1972. ISSN 0036-1399, 1095-712X. doi: 10.1137/0123044. URL http://epubs.siam.org/doi/10.1137/0123044.

[73] Michael Denker, Cristiano Köhler, Aitor Morales-Gregorio, Alexander Kleinjohann, Peter Bouss, Alessandra Stella, Regimantas Jurkus, Maximilian Kramer, Paulina Dabrowska, Moritz Kern, Anno Christopher Kurth, Robin Gutzen, Florian Porrmann, and Sarah Pilz. Elephant 0.11.1, April 2022. URL https://zenodo.org/record/6470226.

[74] P. Du, W. A. Kibbe, and S. M. Lin. Improved peak detection in mass spectrum by incorporating continuous wavelet transform-based pattern matching. *Bioinformatics*, 22(17):2059–2065, September 2006. ISSN 1367-4803, 1460-2059. doi: 10.1093/bioinformatics/btl355.

[75] Sebastian Spreizer, Jessica Mitchell, Jakob Jordan, Willem Wybo, Anno Kurth, Stine Brekke Vennemo, Jari Pronold, Guido Trensch, Mohamed Ayssar Benelhedi, Dennis Terhorst, Jochen Martin Eppler, Håkon Mørk, Charl Linssen, Johanna Senk, Melissa Lober, Abigail Morrison, Steffen Graber, Susanne Kunkel, Robin Gutzen, and Hans Ekkehard Plesser. NEST 3.3, 2022. URL https://zenodo.org/record/6368023.

Extended data figures

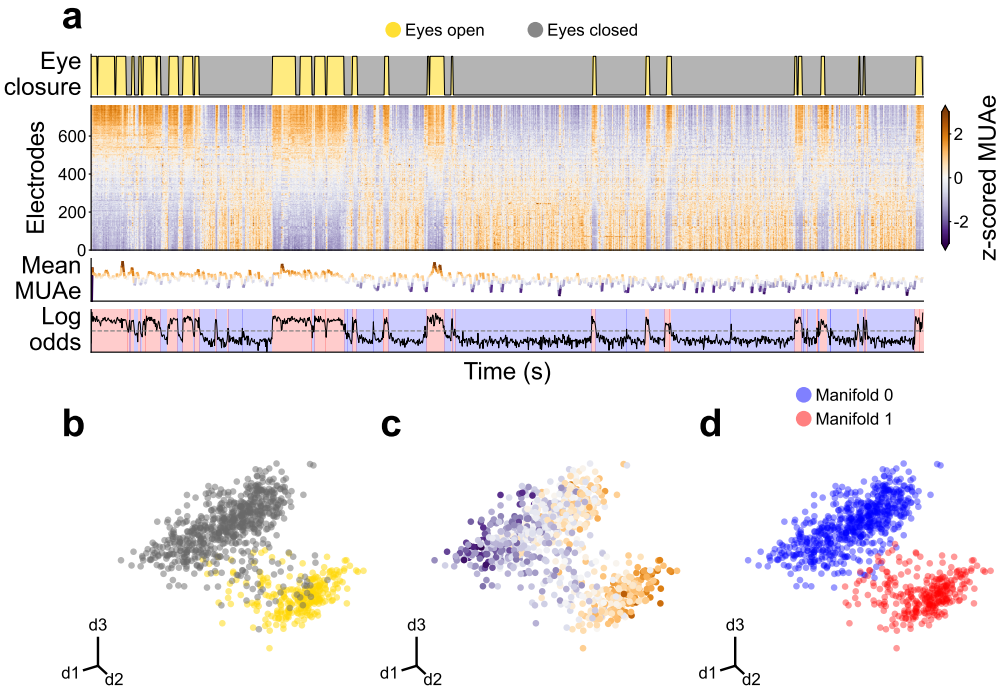

Figure S1: Overview of the experimental data from session L_RS_090817. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

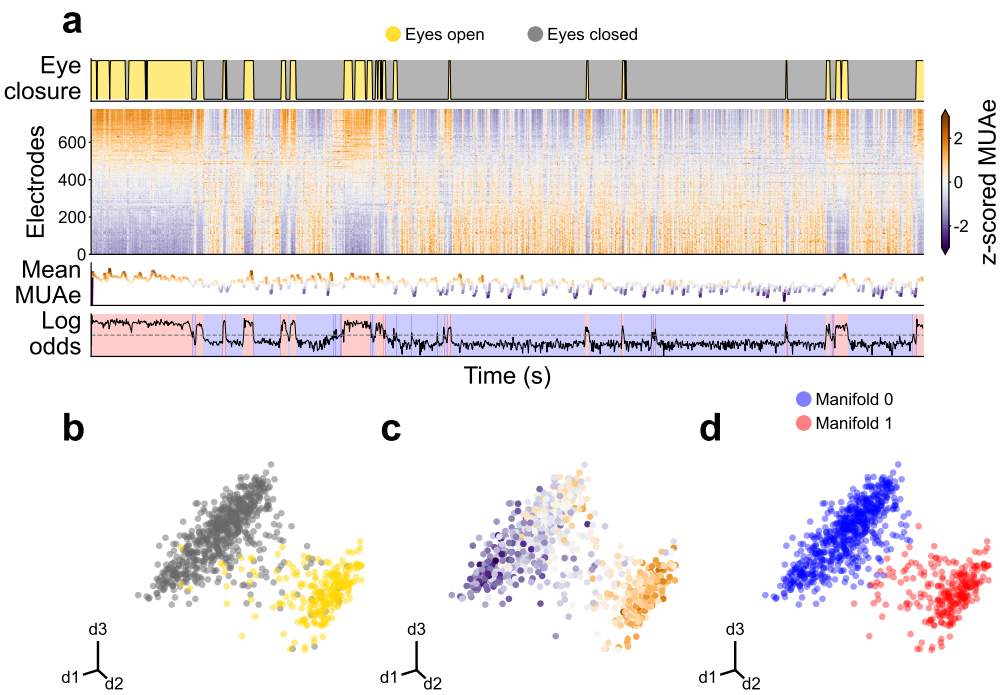

Figure S2: Overview of the experimental data from session L_RS_100817. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

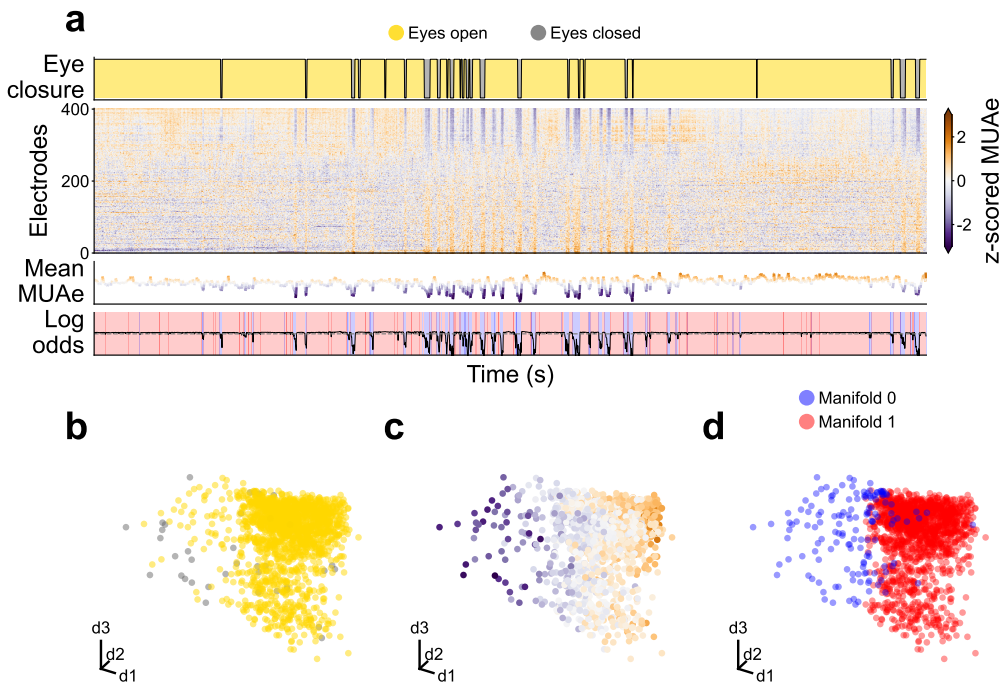

Figure S3: Overview of the experimental data from session A_RS_150819. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

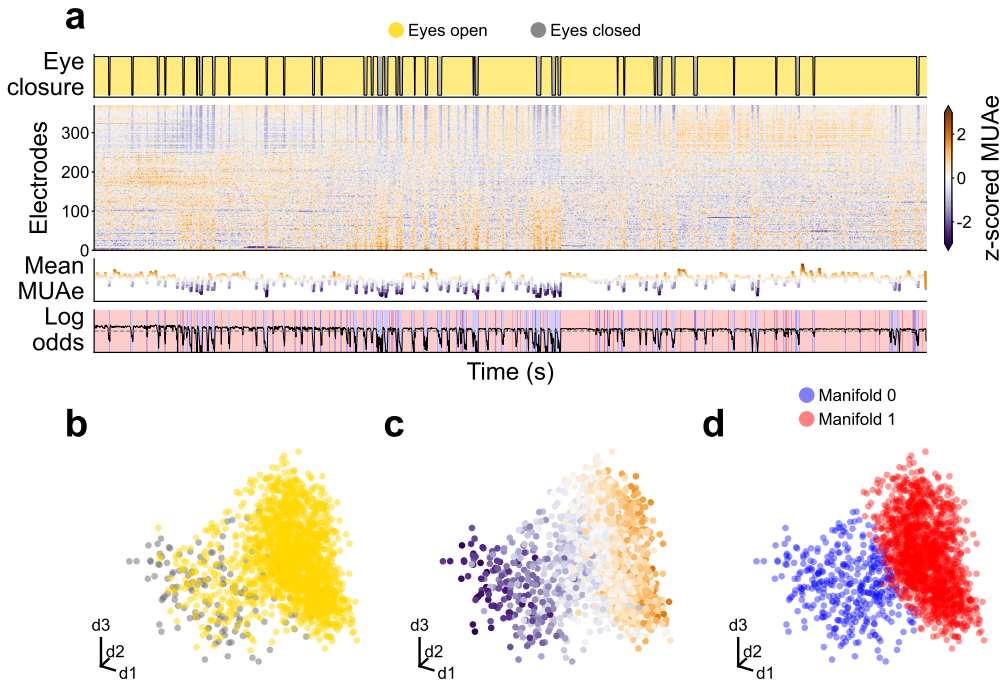

Figure S4: Overview of the experimental data from session A_RS_160819. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

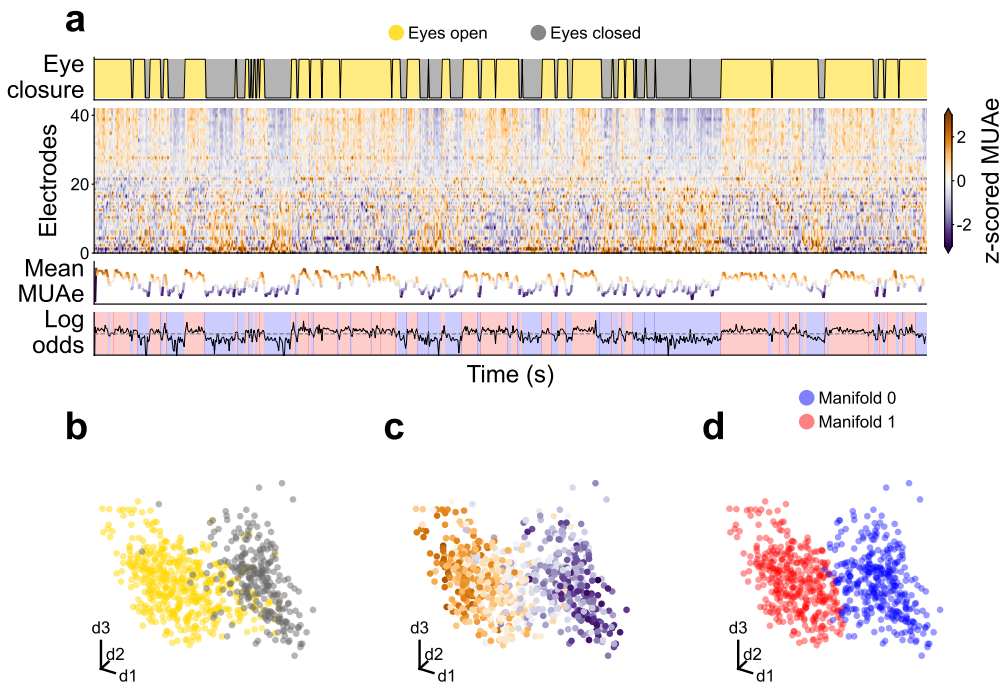

Figure S5: Overview of the experimental data from session Y_RS_180122. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

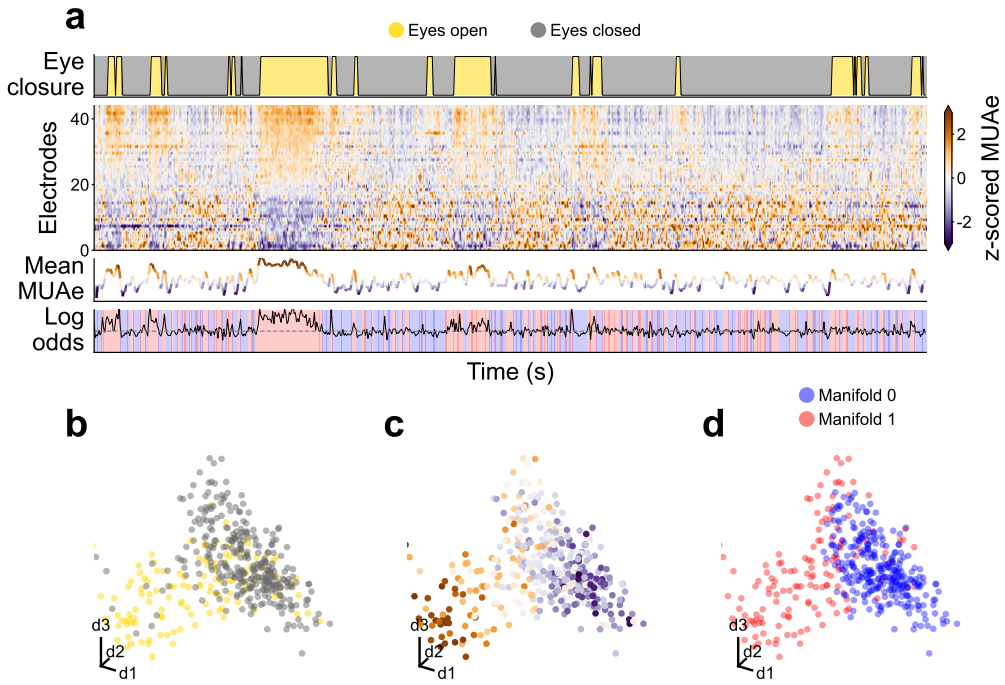

Figure S6: Overview of the experimental data from session Y_RS_180122. **a** Time evolution of signals. **b, c, d** First three principal components of the MUAe neural manifold.

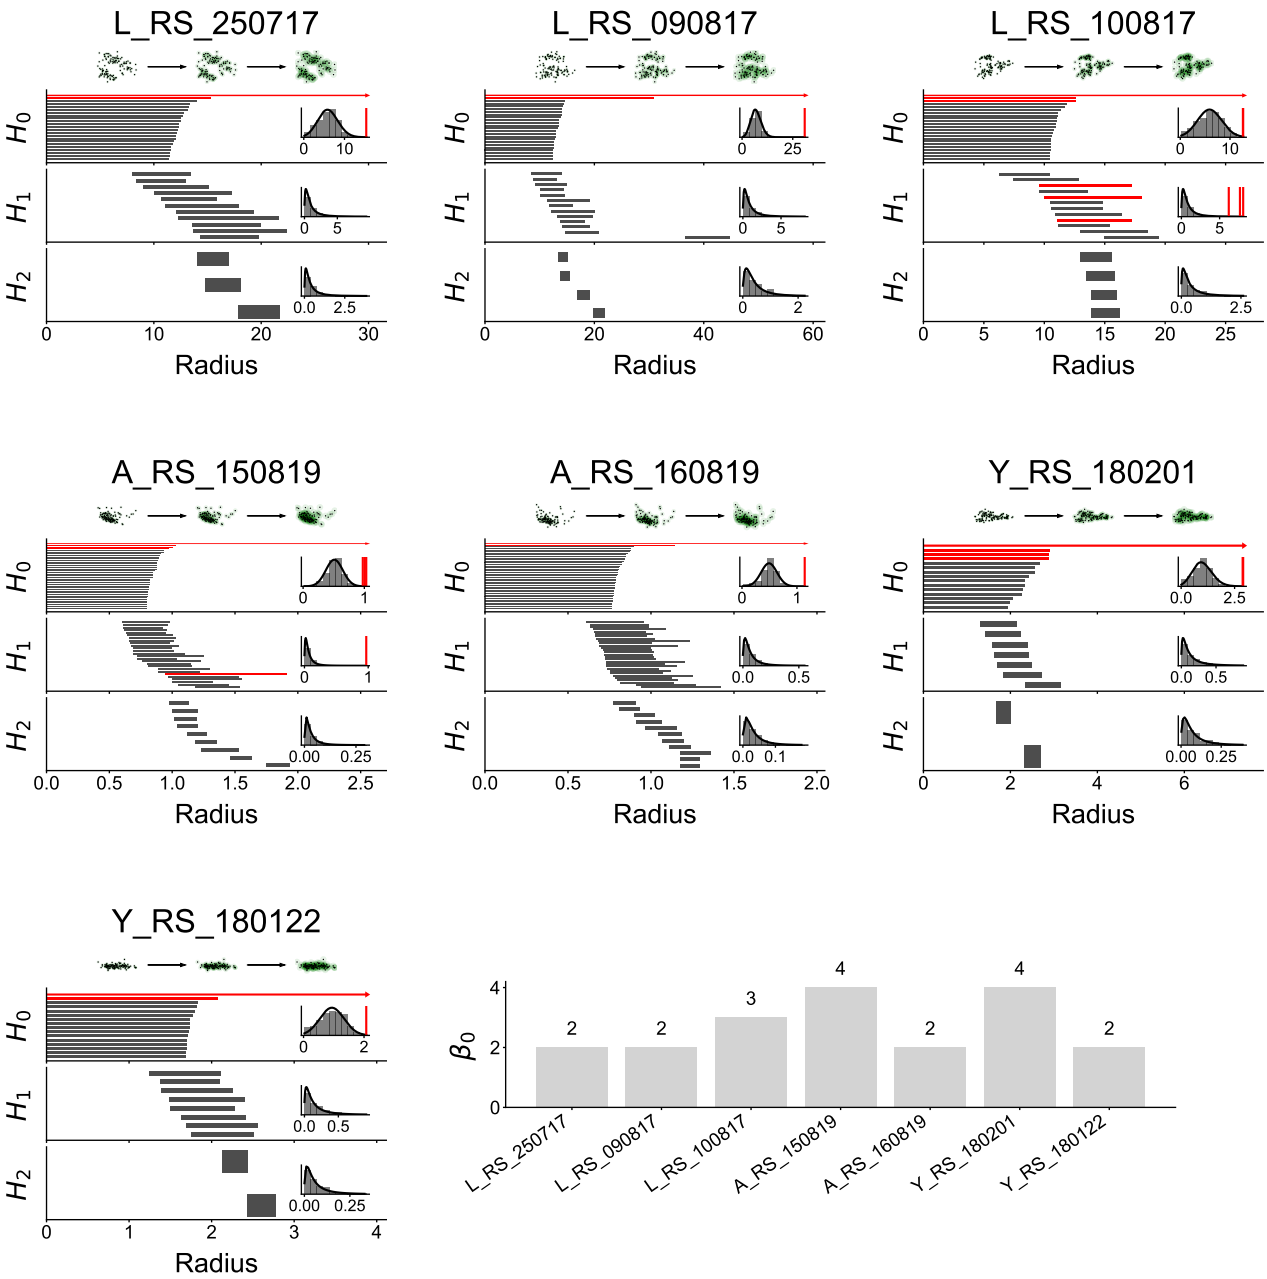

Figure S7: Persistence homology of the high-dimensional manifolds show the presence of at least two clusters. Each panel shows data for one session. For each panel, (Top) Sample clouds with a green radius around them. These correspond to the radius used in the persistent homology process. (Main plots) Persistence barcodes of the Vietoris-Rips complex of the 10D neural manifolds, for all sessions. (Inset plots) Distribution of barcode length with a fitted lognormal distribution, long barcodes coloured red (determined ad hoc). (Bottom right panel) Number of clusters found in each session.

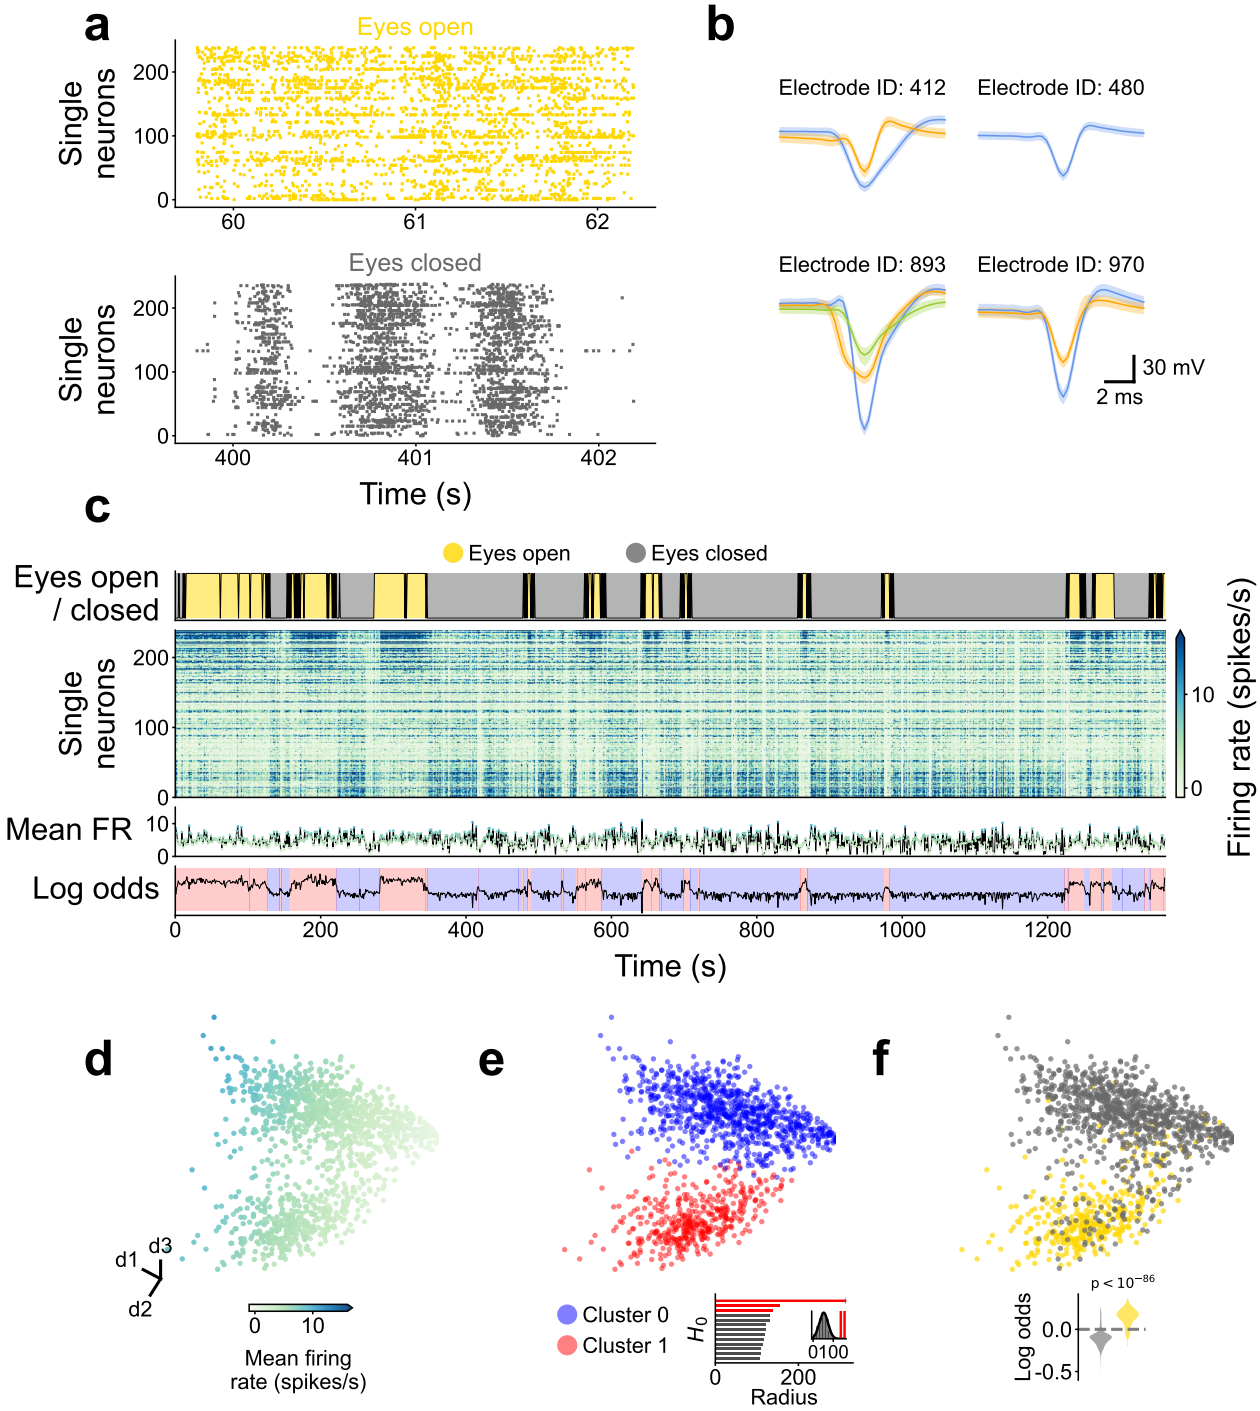

Figure S8: Overview of the spiking data from session L_RS_250717. Single neurons were isolated using a semi-automatic spike sorting method, see Methods, Spike sorting. The firing rate (FR) was calculated counting the number of spikes in 1-second bins. **a** Sample spike raster plots for eyes-open and eyes-closed periods. **b** Sample waveforms from four electrodes, multiple single units isolated in some electrodes (colour-coded). Median (solid line) and 20-80 percentiles (shading) shown per unit. **c** Time evolution of signals. **d, e, f** First three principal components of the firing rate (FR) neural manifold. Insets show the persistent homology for the H_0 homology group (e) and the violin plots of the eye closure against the clustering log odds (f).

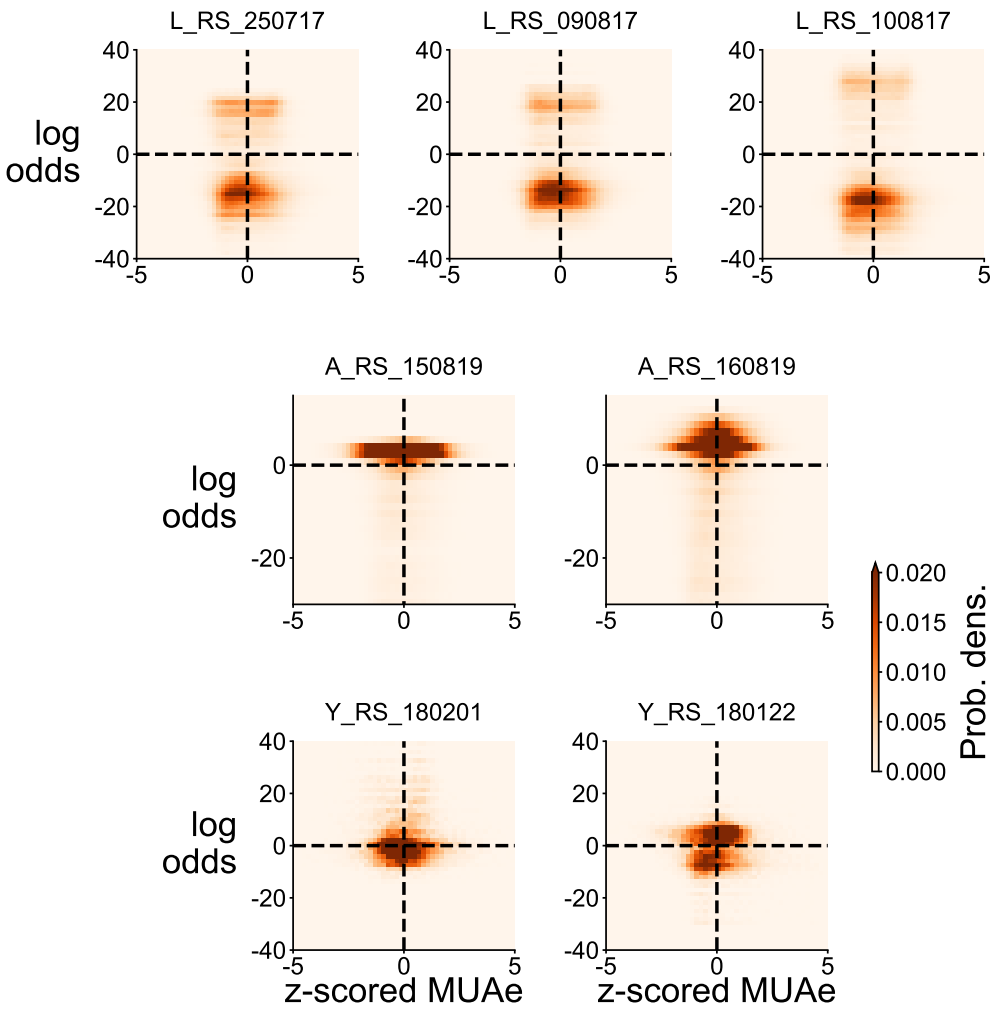

Figure S9: 2D histograms of z-scored MUAe and log odds. Darker colour indicates higher occurrence. If the neural manifolds were solely explained by the higher activity, the histograms should be strictly diagonal; we instead observe that the histograms spread across multiple quadrants and are even bimodal in some cases.

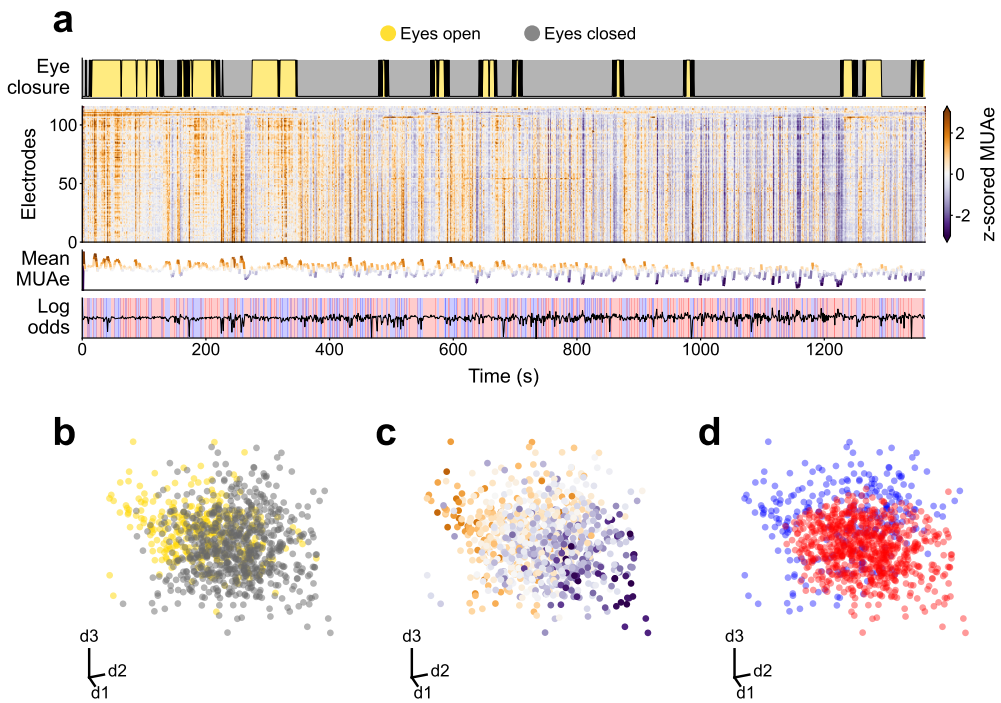

Figure S10: V4 activity from session L_RS_250717 does not show distinct clusters in its neural manifold. **a** Time evolution of signals. **b, c, d** Three dimensional PCA of the MUAe neural manifold.

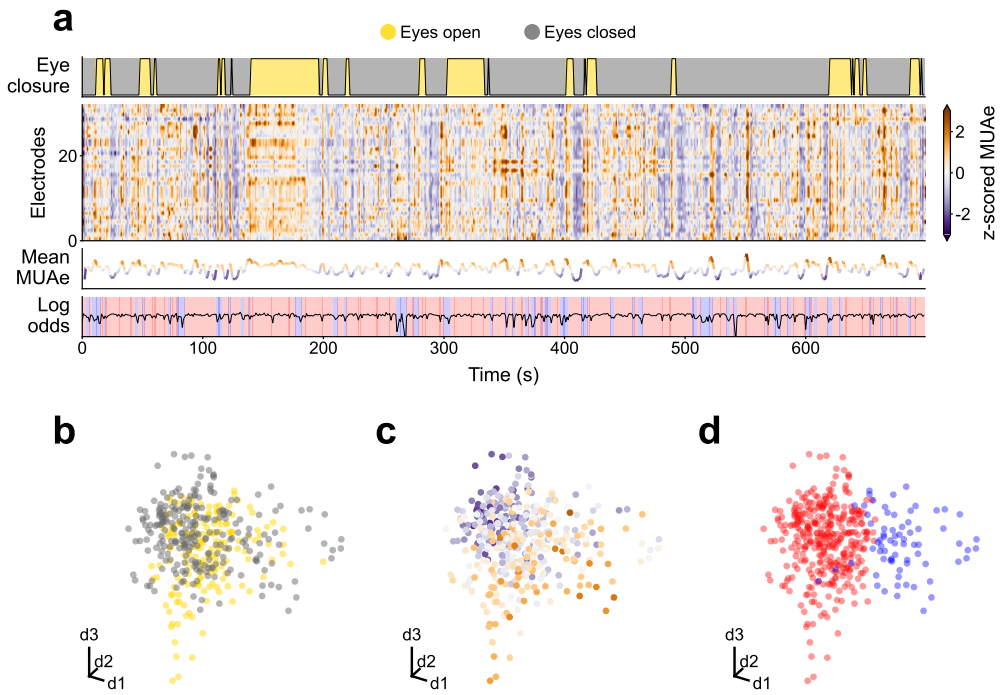

Figure S11: DP activity from session Y_RS_180201 does not show distinct clusters in its neural manifold. **a** Time evolution of signals. **b, c, d** Three dimensional PCA of the MUAe neural manifold.

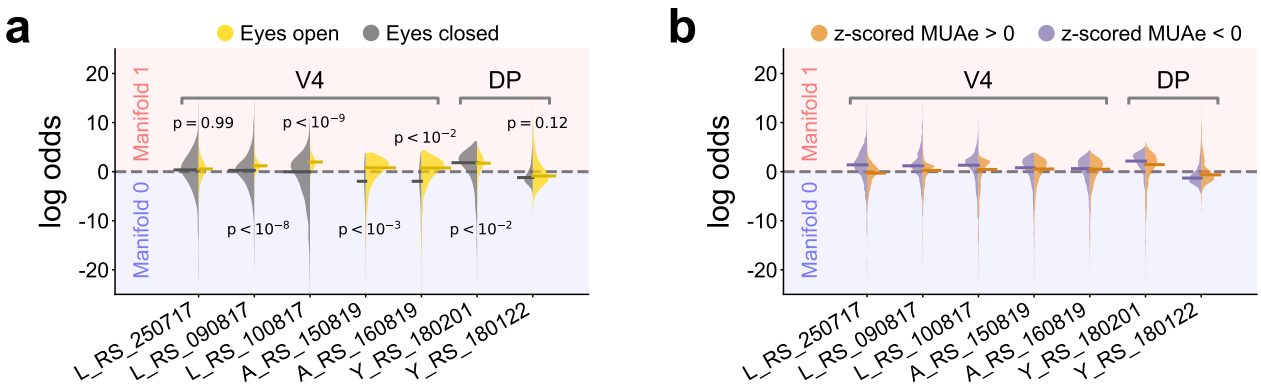

Figure S12: V4 and DP manifold log odds are not strongly correlated with eye closure nor with MUAe. **a, b** Violin plots of V4 and DP for eye closure (**a**) and MUAe activity (**b**). Neither show a clear separation along different neural manifolds.

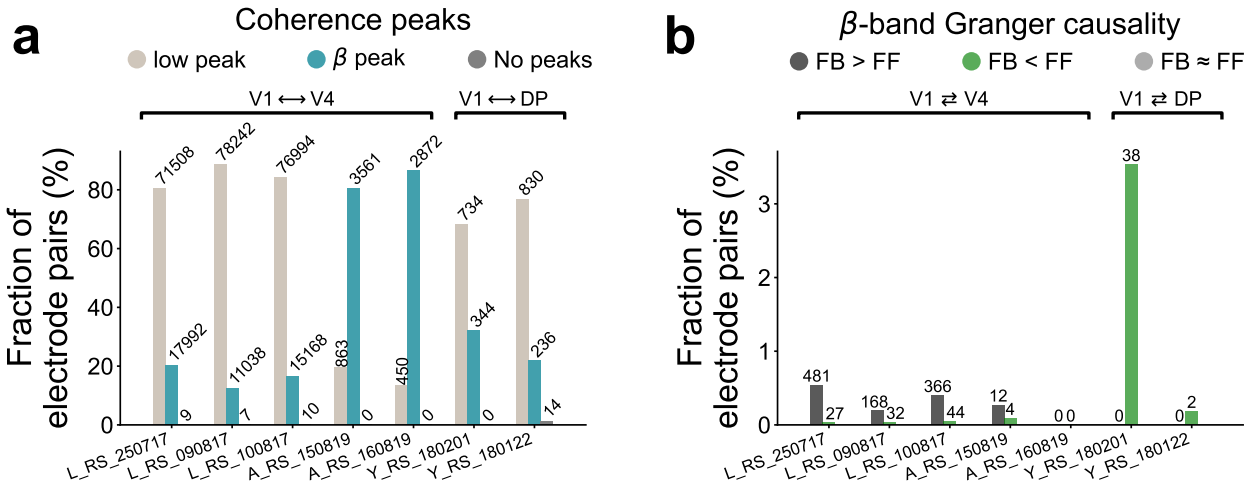

Figure S13: Quantification of coherence peaks and beta-band spectral Granger causality. **a** Quantification of coherence peaks across all sessions. A substantial portion of all electrode pairs displayed a beta peak. Note that the percentages for a session can add up to more than 100% since the same electrode pair can have both a low-frequency and a beta peak. **b** Quantification of beta-band spectral Granger causality for all sessions. Welch's t-test was used to determine whether top-down Granger causality was greater than, less than, or roughly equal to bottom-up Granger causality, within the beta frequency band. The test was only applied to those electrode pairs that showed a beta coherence peak. A large portion of V1 ⇄ V4 pairs show stronger causality in the top-down direction, while V1 ⇄ DP did not appear to have prominent top-down causality compared to bottom-up causality.

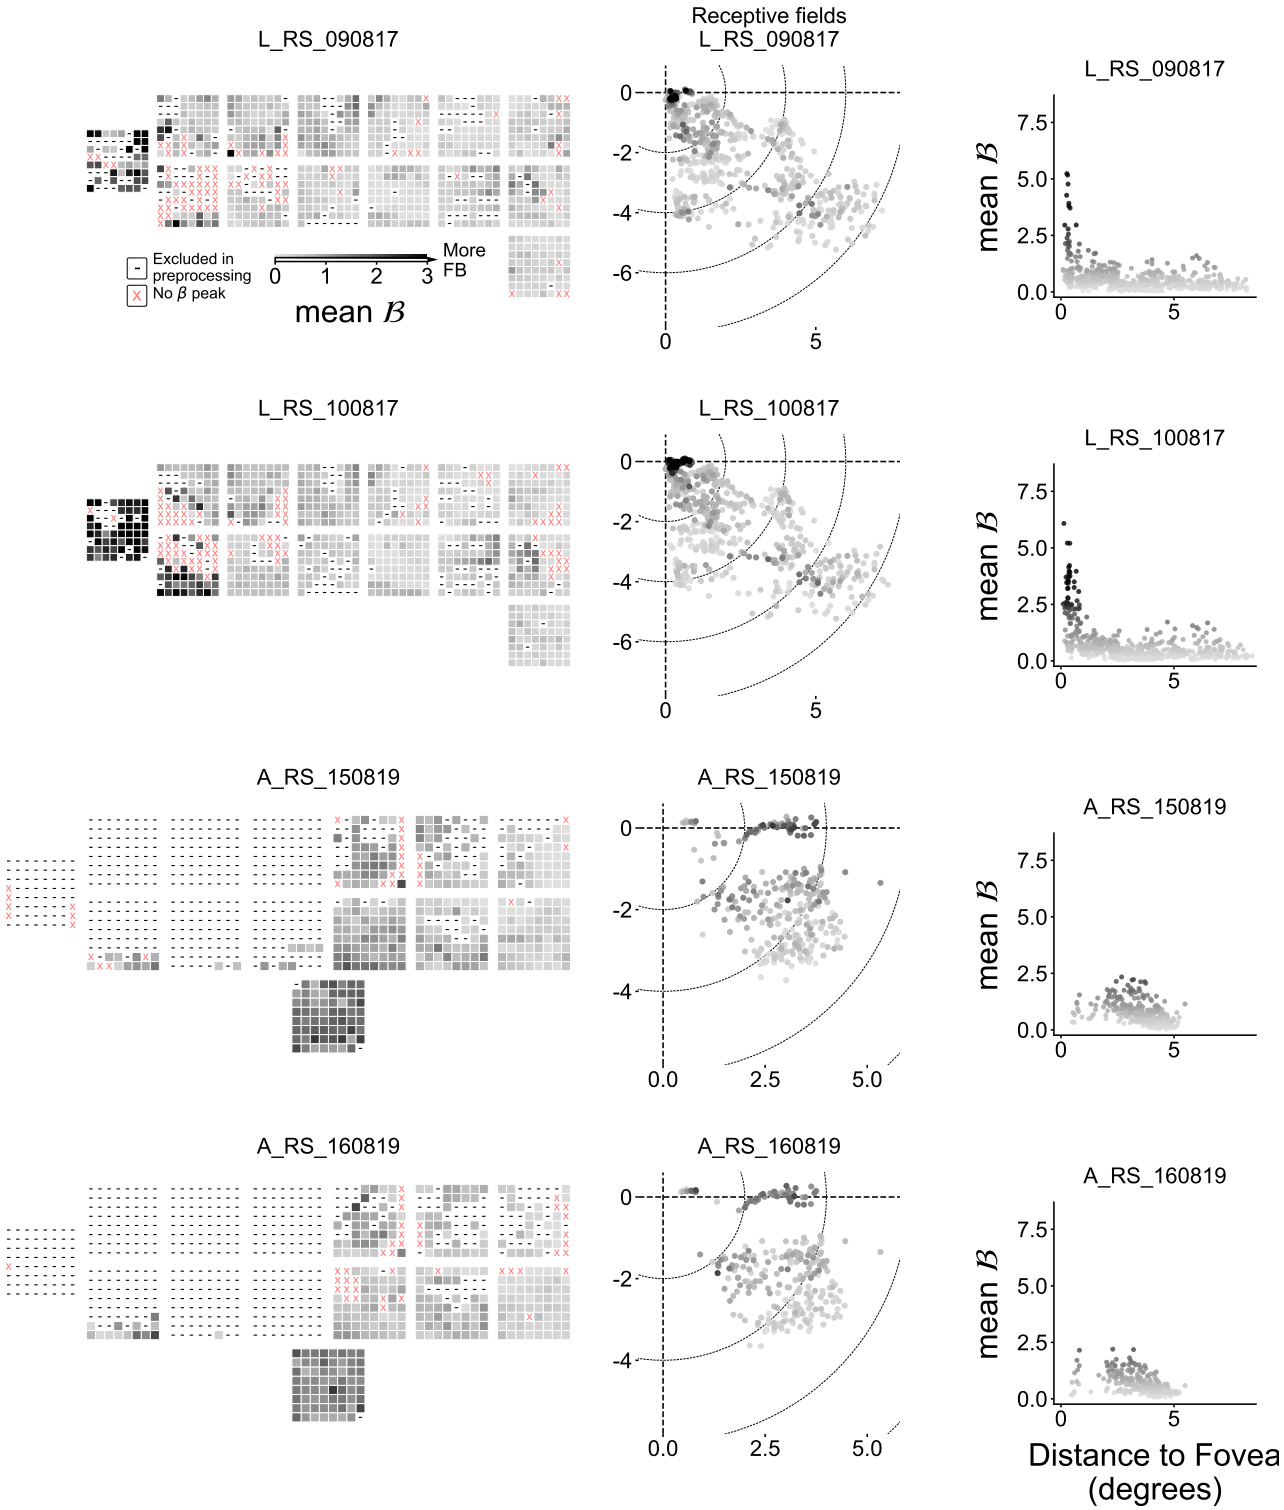

Figure S14: Spatial distribution of Granger causality strength per electrode for all relevant sessions. *(Left column)* Schematic representation of the electrode locations overlaid with the mean top-down signal strength $\mathcal{B}$ per electrode (see Coherence and Granger causality for a description of $\mathcal{B}$). *(Center column)* Receptive field (RF) map overlaid with the mean $\mathcal{B}$ per electrode. Stronger $\mathcal{B}$ is found around the foveal region of V1. *(Right column)* Mean $\mathcal{B}$ displayed against the distance from the fovea.

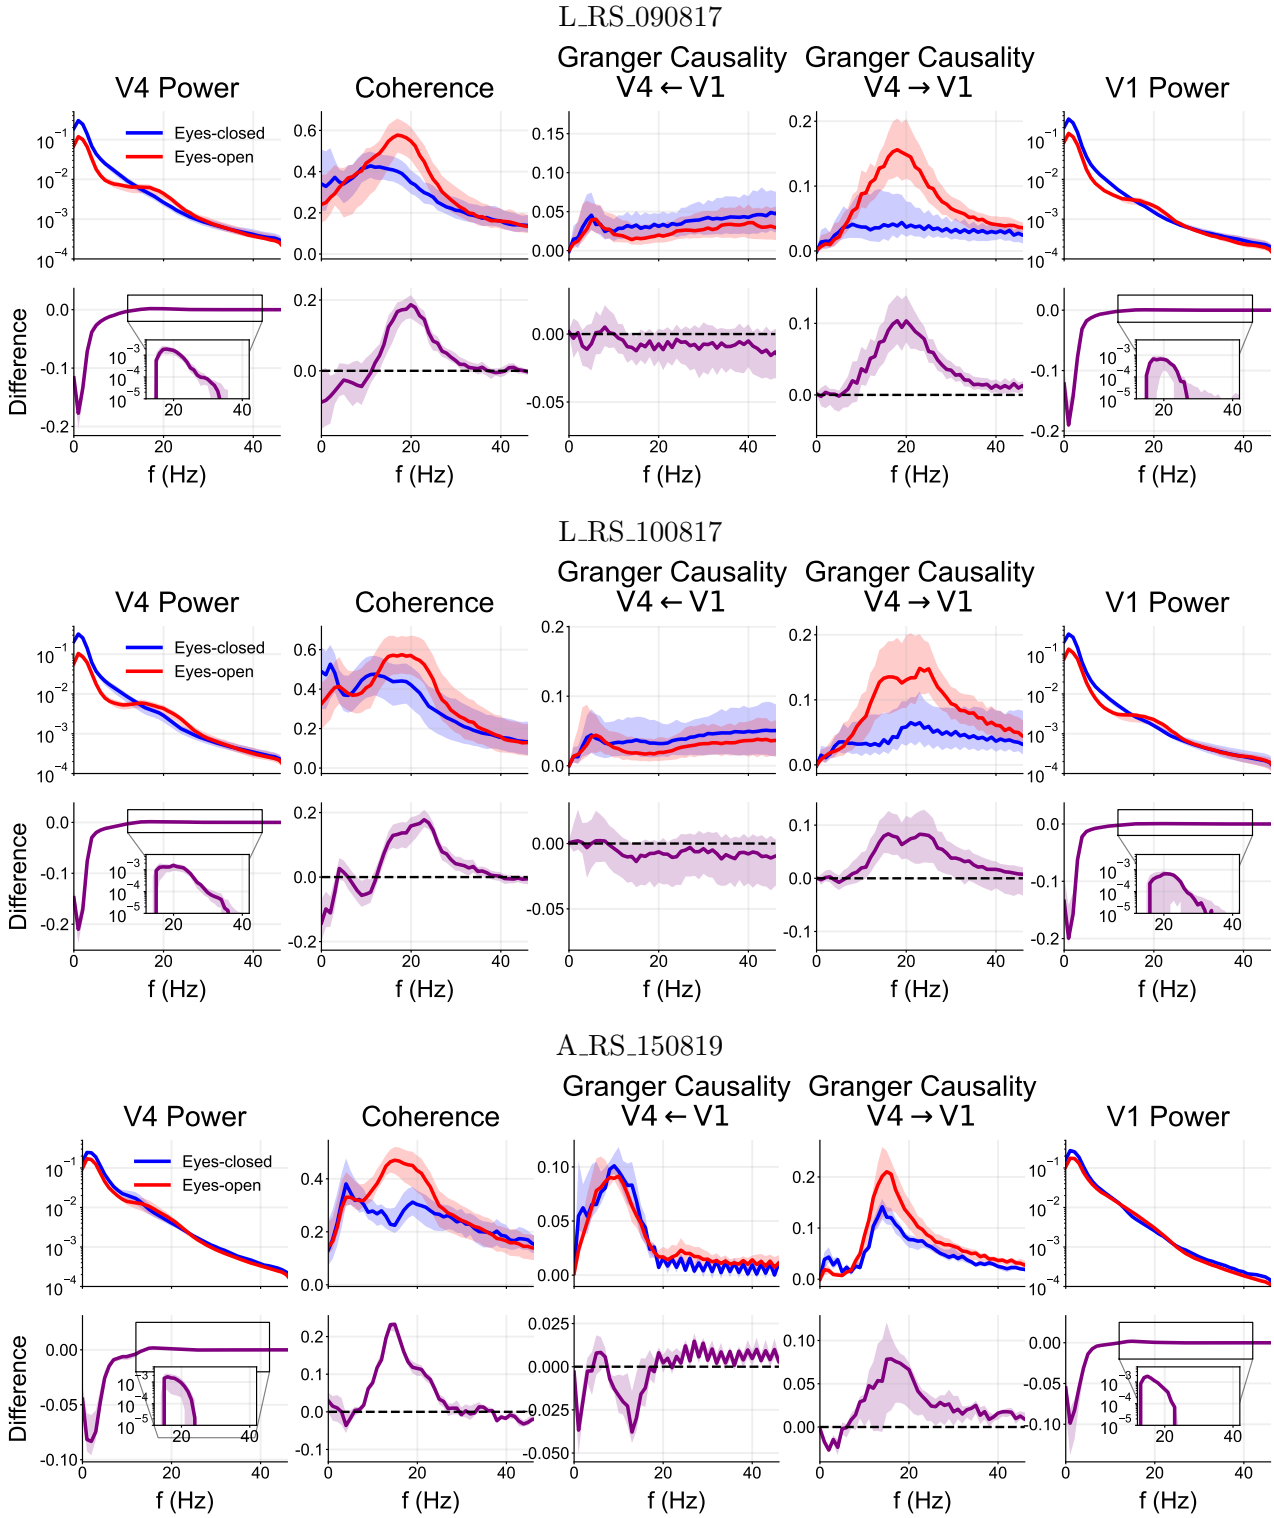

Figure S15: Spectral power, coherence, and Granger causality for the electrodes with high causality strength ($\mathcal{B} > 10$) in sessions L_RS_090817, L_RS_100817, and A_RS_150819. The data for each behavioural condition (eyes-open/closed) were concatenated and their metrics reported separately. Thick line shows median across electrodes (or pairs of electrodes) and shading indicates the 25th to 75th percentile (top row for each session). The difference between eyes-open and eyes-closed was calculated for each electrode or pair of electrodes (bottom row for each session).

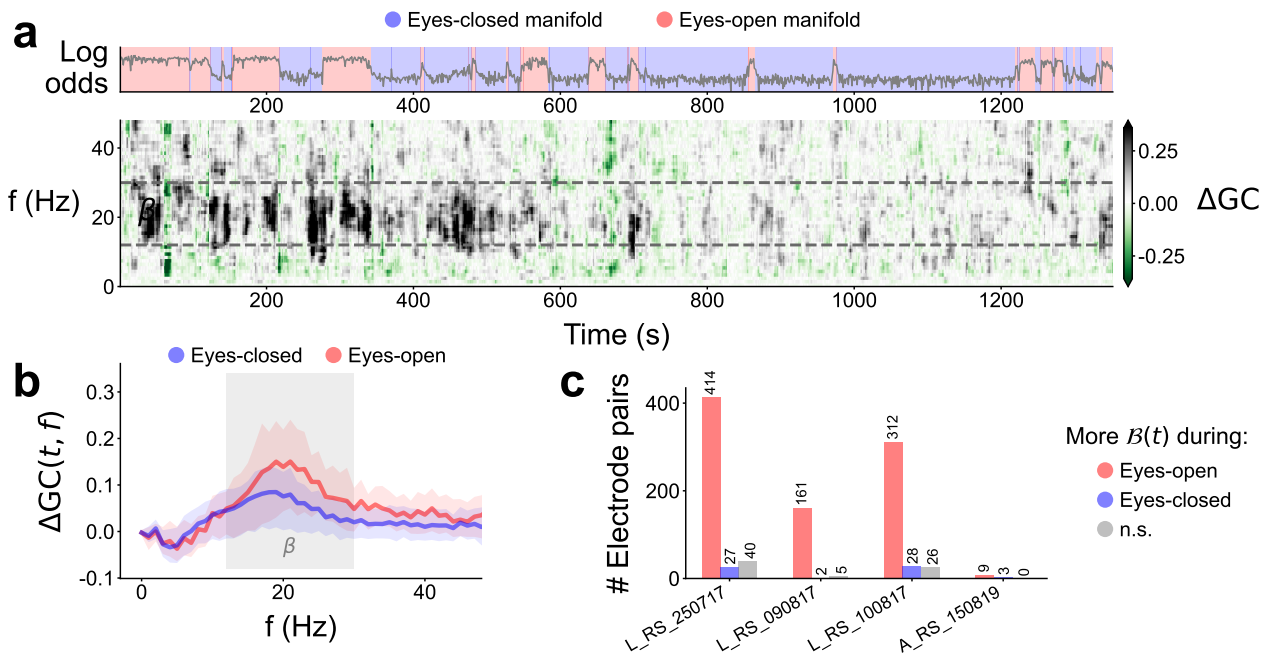

Figure S16: Time-dependent spectral Granger causality reveal higher top-down signals in the eyes-open periods. **a** Time evolution of the log odds (top), and the spectral Granger causality difference for a representative sample of V1-V4 electrodes (bottom). The sample electrodes were the same as in Figure 5. **b** Causality difference median (line) and 25th to 75th percentiles (shade) in each manifold for one sample V1-V4 electrode pair. Beta frequency range highlighted. **c** Quantification of beta-band causality difference $\mathcal{B}(t)$ over time (in each manifold) for all V1-V4 and V1-DP electrodes in all sessions—using Welch's t-test.

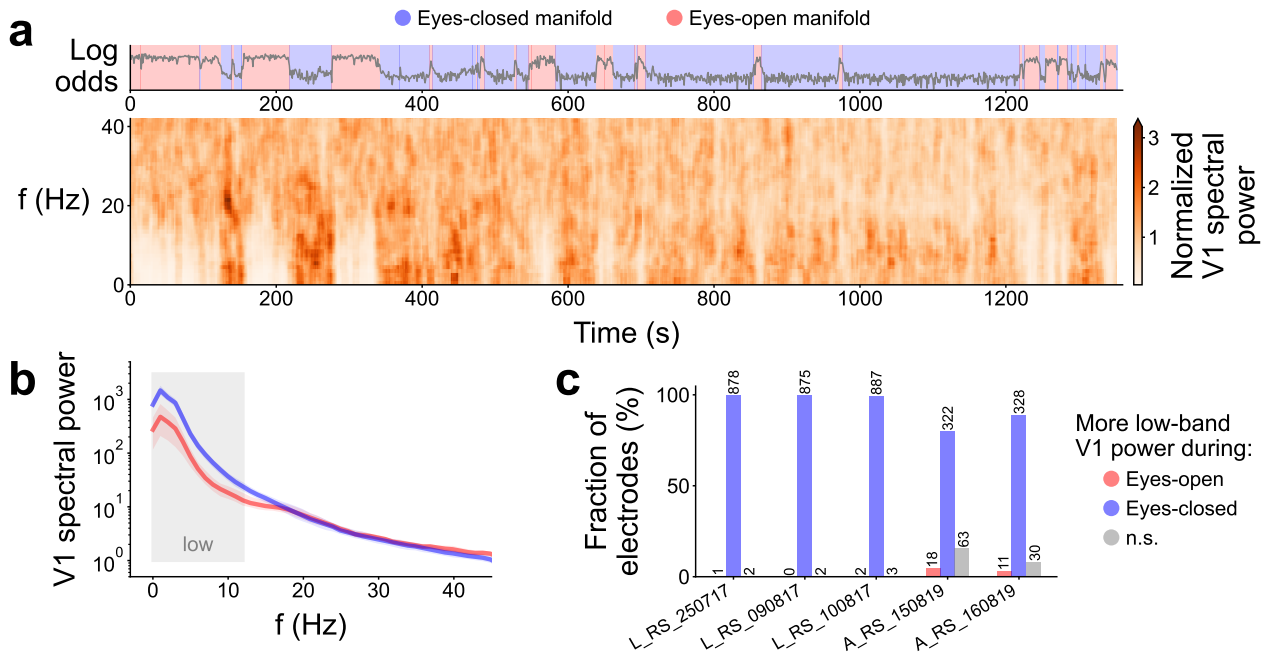

Figure S17: Analysis of the V1 LFP spectrogram. **a** Log odds identifying the neural manifolds (as in Figure 2a), and time-varying spectrum of a sample V1 electrode (session L_RS_250717, power normalised for each frequency). **b** Spectrum of a sample V1 electrode (session L_RS_250717). Colours indicate the different manifolds. **c** Result of t-test in the low frequency band (less than 12 Hz) for all V1 electrodes. As expected, the overwhelming majority of electrodes displays higher low frequency power when the eyes are closed.

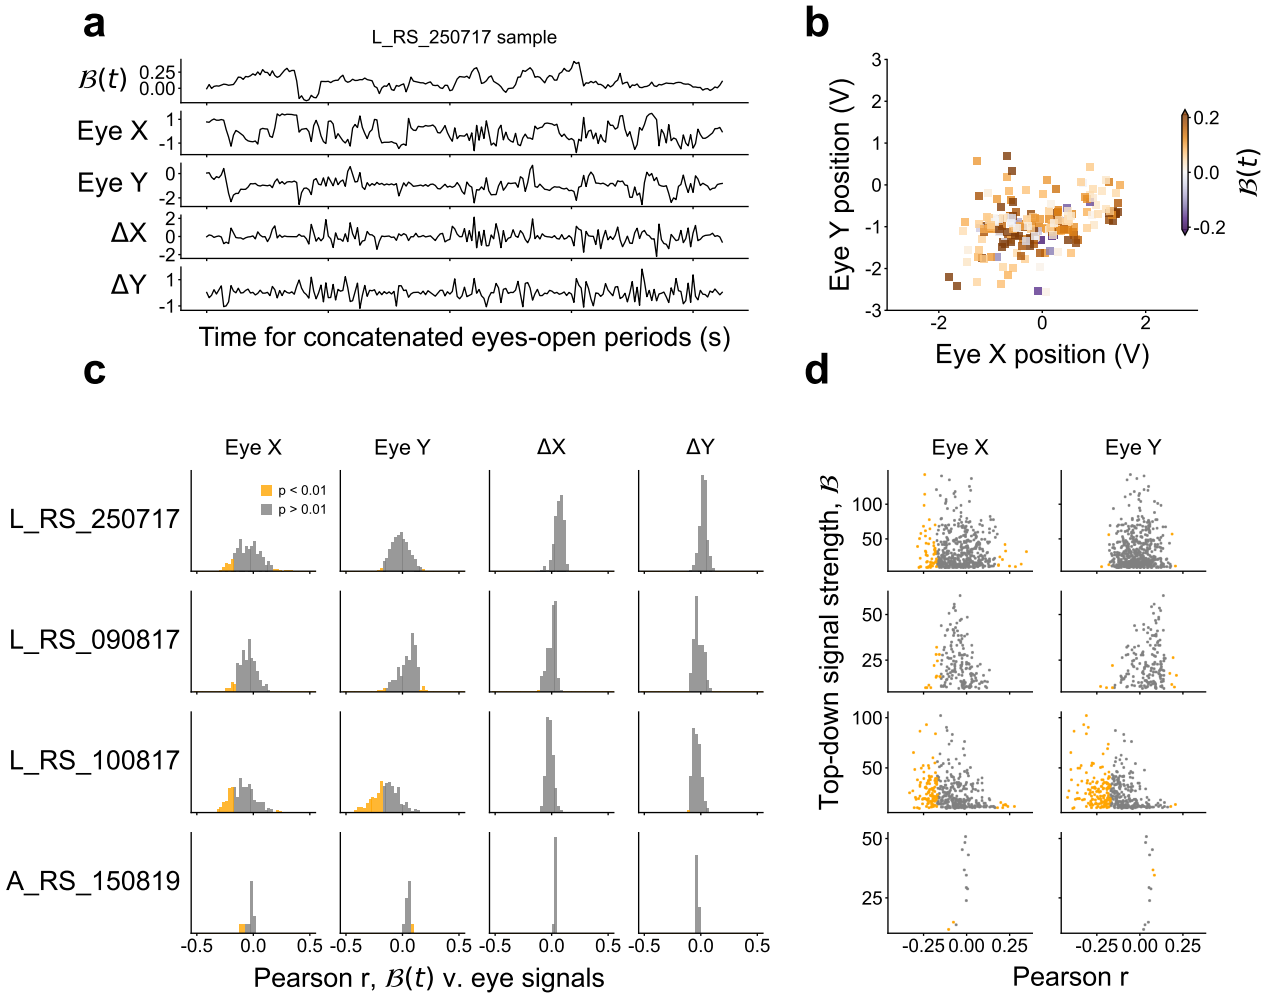

Figure S18: Top-down signals are not correlated with gaze direction. **a** Sample traces of the mean beta-band Granger causality difference, gaze direction (Eye X, Eye Y), and gaze direction derivative (ΔX, ΔY). **b** Sample beta causality difference over the gaze locations. Higher top-down causality is not concentrated in particular regions. **c** Histograms of Pearson correlation coefficients between time-dependent causality difference and gaze signals, computed for all electrode pairs in all sessions. Significant ($p < 0.01$ two-tailed) part of histograms shown in orange. Gaze direction derivatives show no significant correlations. Note that we did not correct for multiple testing, since reducing the p-value threshold would simply reinforce our finding that no strong correlation was present between the gaze and the top-down signals; i.e., no multiple testing correction was the more conservative approach in this case. **d** Scatter plot of the summed time-independent causality difference against the correlation with gaze direction. There is no clear relation between $\mathcal{B}(t)$-gaze correlation and causality strength.

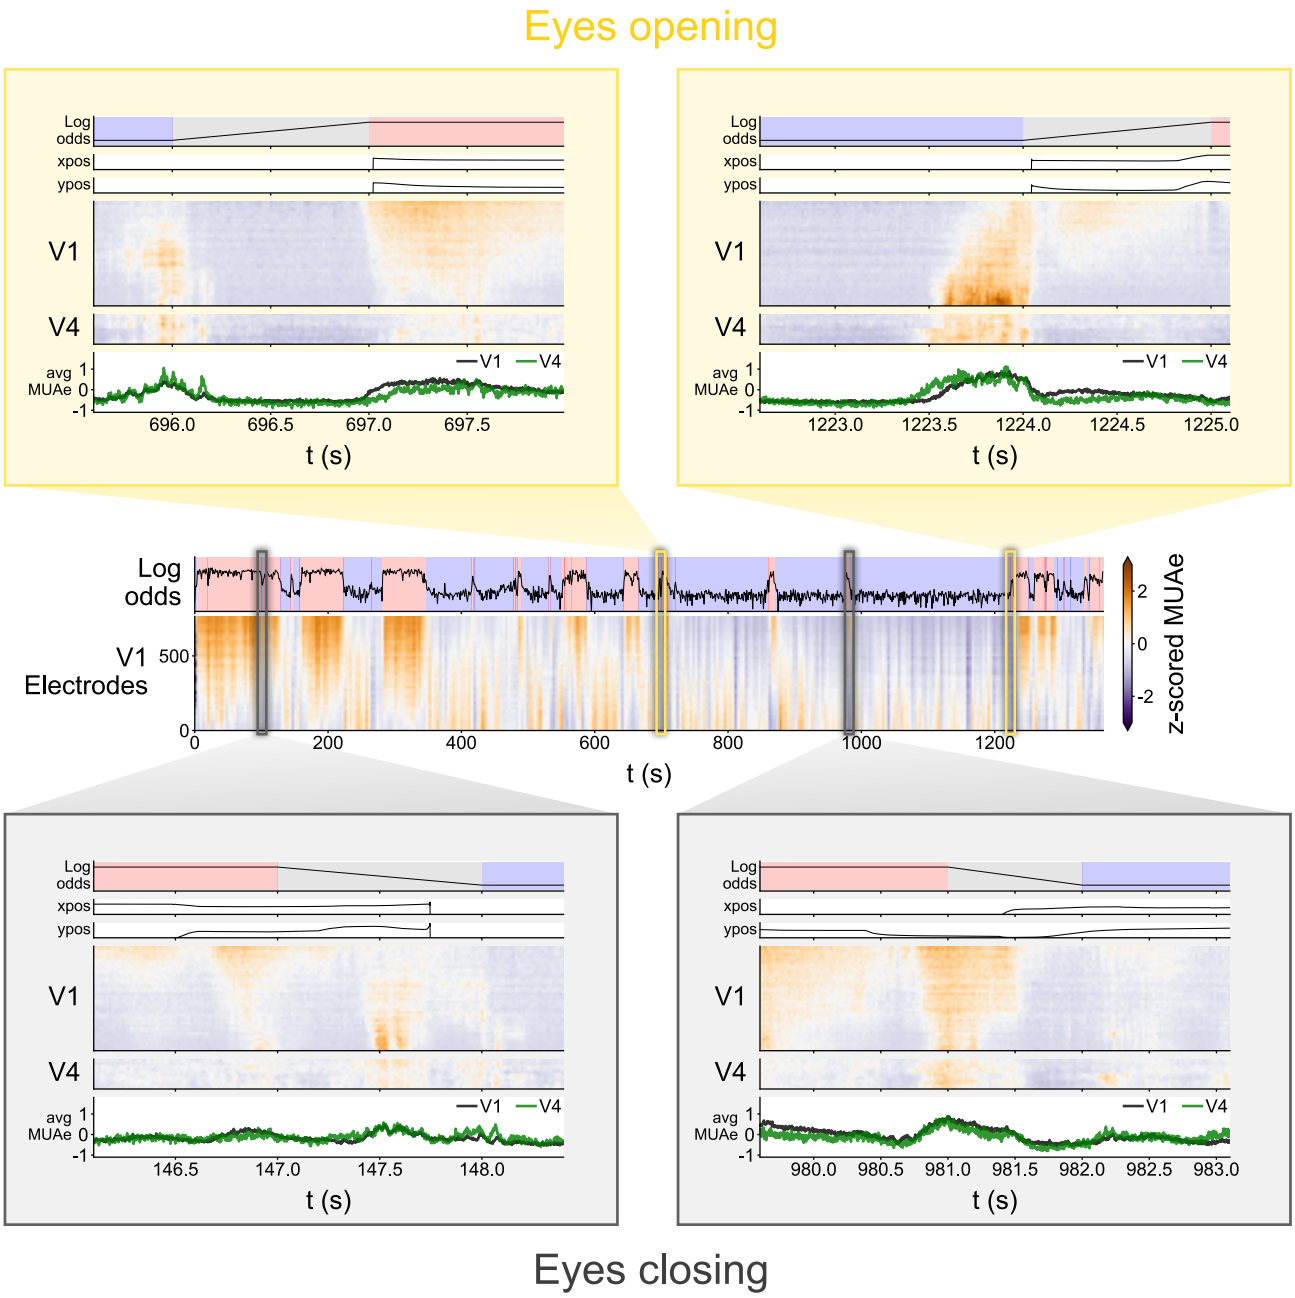

Figure S19: Closer look at the V1 and V4 MUAe around the transitions between the two manifolds in session L_RS_250717.