Fast-local and slow-global neural ensembles in the mouse brain

Thomas J. Delaney[a] **and Cian O'Donnell**[a,b,1]

[a]School of Computer Science, Electrical and Electronic Engineering, and Engineering Mathematics, University of Bristol, Bristol, BS8 1UB, UK; [b]School of Computing, Engineering and Intelligent Systems, Ulster University, Derry/Londonderry, BT48 7JL, UK

Ensembles of neurons are thought to be co-active when participating in brain computations. However, it is unclear what principles determine whether an ensemble remains localised within a single brain region, or spans multiple brain regions. To address this, we analysed electrophysiological neural population data from hundreds of neurons recorded simultaneously across nine brain regions in awake mice. At fast sub-second timescales, spike count correlations between pairs of neurons in the same brain region were stronger than for pairs of neurons spread across different brain regions. In contrast at slower timescales, within- and between-region spike count correlations were similar. Correlations between high-firing-rate neuron pairs showed a stronger dependence on timescale than low-firing-rate neuron pairs. We applied an ensemble detection algorithm to the neural correlation data and found that at fast timescales each ensemble was mostly contained within a single brain region, whereas at slower timescales ensembles spanned multiple brain regions. These results suggest that the mouse brain performs fast-local and slow-global computations in parallel.

neural ensembles | neural correlations | whole-brain computation | neural data science

The brain is traditionally parcellated into anatomical regions that perform distinct computations (1). However these regions do not operate independently: successful brain function must also involve computations spread over multiple regions (2–4). It is unclear how local computations within a single brain region are coordinated with global computations spread across many brain regions. Several possibilities have been proposed: synchronous oscillatory activity may bind together spatially separated neural signals (5–8); travelling waves may propagate signals across the cortex (9); or a hierarchy of timescales may separate low-level sensory processing from higher-level cognitive computations in the brain (10–12).

Here we tested the hypothesis that computations are local to single brain regions at fast timescales, but spread across multiple regions at slower timescales.

Results

Spatial extent of neural correlations varies with timescale. We first characterised the magnitudes of within- and between-region neural spike count correlations by analysing previously published data from ∼500 neurons recorded simultaneously across 9 brain regions (frontal, sensorimotor, visual, and retrosplenial cortex, hippocampus, striatum, thalamus, and midbrain) in awake mice (13, 14). We calculated spike count correlations for each pair of neurons in the dataset over a range of different time bin widths, from 10 milliseconds to 3 seconds.

Figure 1C shows example 10-second raster plots and corresponding spike count time series from a pair of neurons within

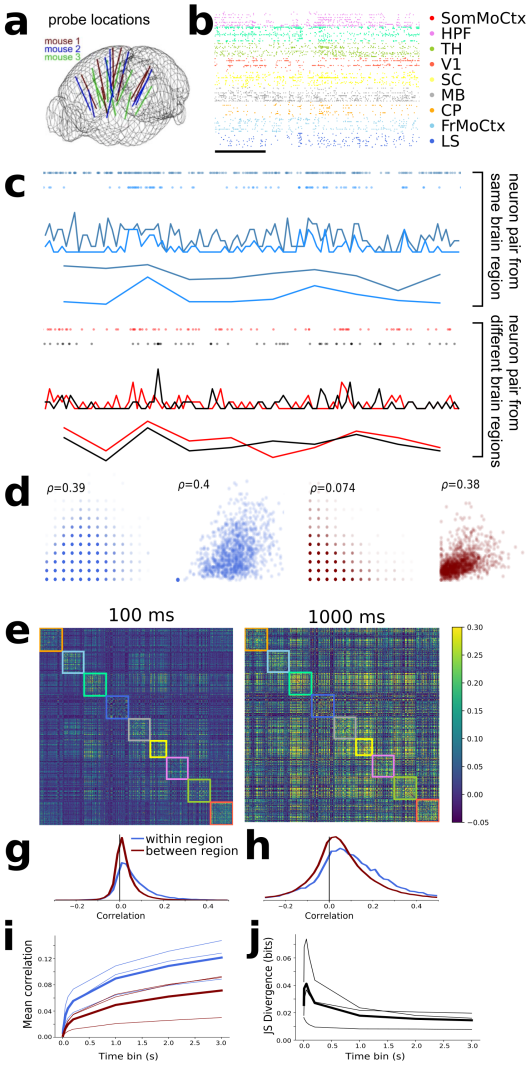

Fig. 1. Within- and between-region neural correlations are more similar at slow timescales than fast timescales. **a:** Neuropixel probe locations in the three mouse brains (adapted from (13)). **b:** Raster plot of spikes from 198 sample units from one mouse. Scale bar corresponds to 1 s. **c:** Spike-count time series from pair of neurons recorded in the same brain region (top) and pair recorded from different regions (bottom). **e, f:** Correlation matrix for spike counts from 494 neurons recorded from one animal with a time bin of 100 ms (c) or 1 second (d). **g, h:** Histograms of pairwise correlations from matrices in c and d for within- and between-region pairs of neurons (colours blue and red respectively) for 100 ms (g) or 1 second (h) time bins. **i:** Mean pairwise correlations as a function of time bin. **j:** Jensen-Shannon divergence of within vs between-region correlation distributions as a function of time bin.

the same brain region (light and dark blue, top) and a pair of neurons from two different brain regions (red and black, bottom) for both 100 ms and 1 s time bins. Figure 1D shows scatter plots of the spike counts for the same neuron pairs. The within-region cell pair showed the same high spike count correlation of $\rho \approx 0.4$ at both 100 ms and 1 s time bins. In contrast, the between-region pair showed a low spike count correlation of 0.07 at fast 100 ms time bins, but a high correlation of 0.4 at slower 1 second time bins. This general pattern held up across the dataset: Figure 1E shows the pairwise correlation matrices for all 494 neurons analysed from this animal for both the 100 ms and 1 second time bin sizes. The rows and columns of these matrices are ordered by brain region, so within-region correlations are inside the coloured boxes along the main diagonal (each colour represents a different brain region). With 100 ms bins, the within-region correlations appear stronger than the between-region correlations. However with 1 second time bins, the within- and between-region correlations appear visually similar. To explore this phenomenon, we separately histogrammed the within- and between-region values from the correlation matrices (Figure 1G,H). Both the mean (Figure 1I) and the width of correlation histograms increased with time bin size, for both within- and between-region correlations (15). However, the within-region correlations had a heavier positive tail than the between-region correlations at fast timescales, but markedly less so at slow timescales (Figure 1G,H). To quantify this effect, we calculated the Jensen-Shannon (JS) divergence between the two distributions. High divergence values imply greater differences in the distributions. Indeed the JS divergence decreased as a function of time bin size, consistently for the data from all three animals (Figure 1J). These results imply that at fast timescales, correlations are high only between neurons within brain regions, but at slow timescales within- and between-region neural correlations are similar.

Low firing rate neurons preferentially correlate within brain region. Low- and high-firing rate neurons have previously been shown to serve different functions in neural circuits (16). To test whether this dissociation is also visible in the within- vs between-region correlation structure, we plotted correlation values against geometric mean firing rate for each pair of neurons in the dataset (Figure 2a–d). Most pairs of neurons had geometric mean firing rates between 1–10 Hz (Figure 2e). Correlations tended to get stronger as a function of firing rate, for both within- and between-region pairs (Figure 2a–d) (17). We binned pairs by their geometric mean firing rate and calculated the JS divergence between the within- and between-region correlations as a function of firing rate bin (Figure 2f). At both fast and slow timescales, low-firing rate pairs had stronger within-region correlations than between-region correlations. In contrast, high firing rate pairs had moderate divergence at 100 ms timebins and almost zero divergence at 1 second time bins. This implies that high-firing rate neurons correlate almost equally strongly within- and between-regions, but low-firing rate pairs have similarly low within- and between-region correlations at all timescales. Therefore

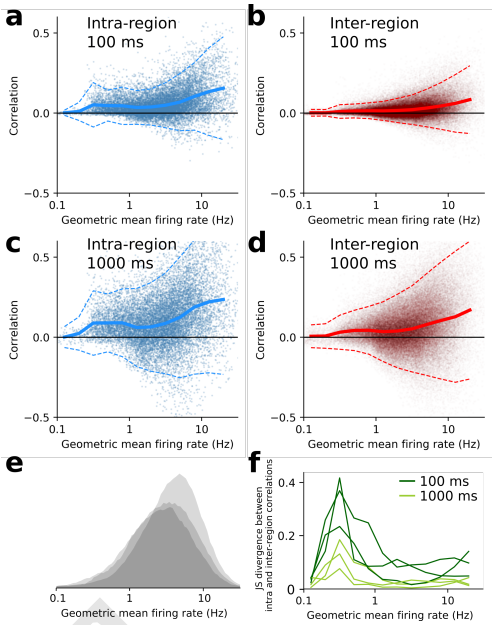

Fig. 2. Low firing rate neurons preferentially correlate within brain regions. **a–d:** Pairwise neural correlations vs geometric mean of firing rate for many pairs from one animal, for within (a,c) and between (b,d) region neuron pairs, with time bin interval shown in panel insets. Solid line shows mean correlation, dashed lines are $\pm$ 2 s.d. from mean. **e:** Histograms of all pairwise geometric mean firing rates for all three animals. **f:** Jensen-Shannon divergence between within and between-region correlations as a function of geometric mean firing rate, for all three animals. Dark green corresponds to spikes binned at 100 ms intervals, light green is 1000 ms intervals.

the phenomenon seen in Figure 1 is mainly due to high-firing rate neuron pairs.

Detected ensembles align with anatomical regions at short time bins, but not long time bins. To test if neural ensembles also showed different structure at fast and slow timescales, we ran a community detection algorithm from network science on the correlation matrices to detect ensembles (Figure 3a) (18). The algorithm splits the neurons into non-overlapping subsets based on their correlations, trying to discover ensembles of neurons with strong positive correlations between the members of each ensemble, but weaker correlations with neurons in other ensembles (Methods). Figure 3b and c shows the same example correlation matrices from Figure 1e, but with the rows and columns reordered by ensemble membership. In all three animals we found fewer ensembles at longer time bin sizes (Figure 3f). Crucially, the ensemble detection algorithm did not know anything about which brain regions each neuron belonged to. To visualise the brain region membership of each ensemble, we plotted a small square for each neuron coloured according to its brain region (Figure 3d,e). At 100 ms time bins, most ensembles contained neurons from only a small number of brain regions, whereas at 1 second time bins almost all ensembles contained neurons from several brain regions. To quantify this effect, we asked the questions: what is the probability that any arbitrary neuron pair is in the same ensemble? And does this differ for pairs of neurons within the same brain region vs pairs across two brain regions? 20–30% of same-region pairs were in the same ensemble, but only 10–20% of different-region pairs were in the same ensemble (Figure 3g). The difference between these two fractions decreased towards

Conceived research: TJD. Performed research: TJD and COD. Wrote the paper: TJD and COD.

The authors have no competing interests to declare.

[1]To whom correspondence should be addressed. E-mail: c.odonnell2@ulster.ac.uk

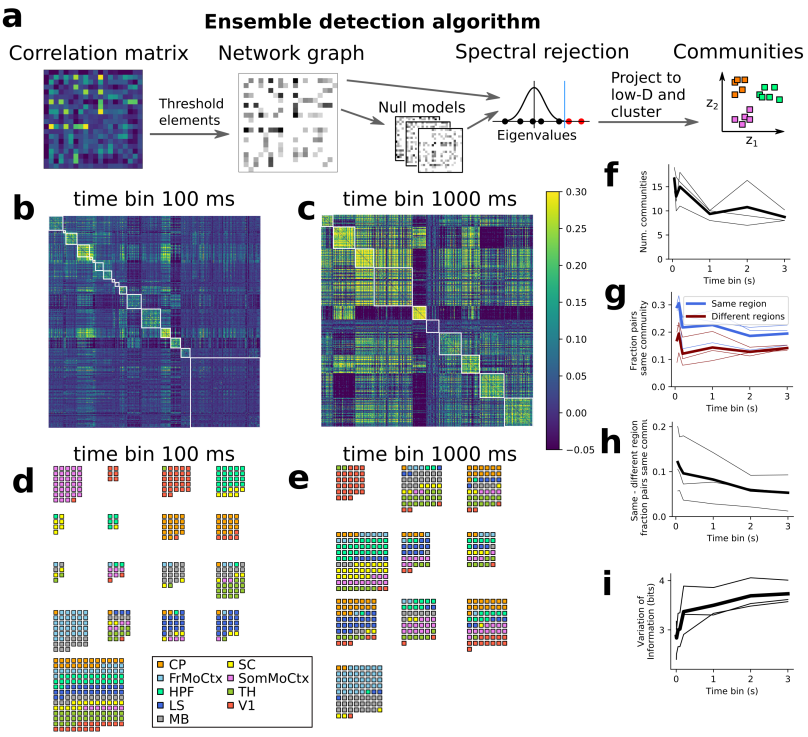

Fig. 3. Neural ensembles are within-region at fast timescales but multi-region at slow timescales. **a:** Schematic diagram of community detection algorithm steps. **b, c:** Same correlation matrices as fig 1 sorted by ensemble. **d, e:** Example ensembles at short (d) and long (e) timescales. **f:** Number of detected ensembles vs time bin size. **g:** Fraction of same and different region neuron pairs being in same community, vs time bin size. **h:** Difference in fraction of same and different region neuron pairs being in same community (same data as panel g). **i:** Variation of information (measure of dissimilarity of anatomical vs ensemble partitions) vs time bin size.

zero as a function of time bin size (Figure 3h), implying that at fast time scales neurons in the same brain region had a higher chance of being in the same ensemble than two neurons in different brain regions, but this distinction got weaker at slower time scales. To further quantify the effect, we used a distance measure from information theory to ask the question: what is the difference between the brain region partition and ensemble partition of the set of all recorded neurons? This 'variation of information' measure increased as a function of time bin size in all three animals (Figure 3i), again implying that anatomical regions and neural activity ensembles are more similar at fast timescales than slow timescales.

Discussion

Although previous studies have compared within- and between-region neural correlations, to our knowledge none have described the fast-local vs slow-global ensembles phenomenon we presented here. There are a few possible reasons for this gap: most electrophysiological studies either looked at small numbers of neurons where the phenomenon may not be statistically detectable, or looked at aggregate neural activity measures such as local field potentials (8) which would miss the single-neuron-resolution ensembles we discovered. Modern large-scale two-photon imaging methods do enable simultaneous recordings from single neurons in multiple brain regions, but with poorer signal-to-noise and slow sampling rate so also may not be able detect the phenomenon we described.

Why might the fast-local vs slow-global dissociation exist? From a mechanistic point of view, one explanation may be that the energetic and space constraints on brain wiring imply that long-range, between-region signals can be transmitted only at low-bandwidth and with some latency (19). There are typically fewer long-range synaptic connections than local connections, between-region signalling is low-dimensional (20), and mammalian axons transmit action potentials between brain regions with latencies of 10–100 ms (21). These bandwidth and latency constraints will limit the speed of any computations that require back-and-forth recurrent signalling between neurons. This issue is well known in human-made computers, where the 'von Neumann bottleneck' for transferring data between memory and CPU via low-bandwidth and high-latency databuses constrains computation speed (22). From a functional point of view, a separation of timescales between local and global computations may allow for less interference between processes, and allow local neural circuits to complete their tasks quickly before broadcasting the results to other brain regions (8).

We examined this phenomenon only for 9 particular brain regions, which despite all exhibiting the effect, differed in their mean firing rates and correlations (13). It would be interesting to try to understand if and how each brain region adapts variations of the general fast-local, slow-global principle depending on its activity statistics and computational role in the brain at large.

Materials and Methods

All data analysed were sourced from a publicly available dataset (14) with experimental procedures described previously (13). Briefly, eight Neuropixel probes were used to record electrophysiological activity simultaneously from nine brain areas: frontal, sensorimotor, visual, and retrosplenial cortex, hippocampus, striatum, thalamus, and midbrain, in each of three 10–16 week-old mice. The mice were awake but head-fixed. We wrote Python code to compute the correlation matrices and implement the community detection algorithm (18). Further details are provided in the Supporting Information.

ACKNOWLEDGMENTS. We thank Michael Ashby for useful discussions and Nick Steinmetz for making the data public. Funding was provided by the EPSRC (Doctoral Training Partnership studentship to TJD) and MRC (MR/S026630/1 to COD).

1. K Zilles, K Amunts, Centenary of brodmann's map—conception and fate. *Nature Reviews Neuroscience* **11**, 139–145 (2010).
2. DS Bassett, ET Bullmore, Human brain networks in health and disease. *Current opinion neurology* **22**, 340 (2009).
3. JD Power, et al., Functional network organization of the human brain. *Neuron* **72**, 665–678 (2011).
4. O Sporns, RF Betzel, Modular brain networks. *Annual review psychology* **67**, 613–640 (2016).
5. CM Gray, P König, AK Engel, W Singer, Oscillatory responses in cat visual cortex exhibit inter-columnar synchronization which reflects global stimulus properties. *Nature* **338**, 334–337 (1989).
6. M Siegel, TH Donner, AK Engel, Spectral fingerprints of large-scale neuronal interactions. *Nature Reviews Neuroscience* **13**, 121–134 (2012).
7. P Fries, Rhythms for cognition: communication through coherence. *Neuron* **88**, 220–235 (2015).
8. TA Engel, ML Schölvinck, CM Lewis, The diversity and specificity of functional connectivity across spatial and temporal scales. *Neuroimage* **245**, 118692 (2021).
9. L Muller, F Chavane, J Reynolds, TJ Sejnowski, Cortical travelling waves: mechanisms and computational principles. *Nature Reviews Neuroscience* **19**, 255–268 (2018).
10. JD Murray, et al., A hierarchy of intrinsic timescales across primate cortex. *Nature neuroscience* **17**, 1661–1663 (2014).
11. JH Siegle, et al., Survey of spiking in the mouse visual system reveals functional hierarchy. *Nature* **592**, 86–92 (2021).
12. R Zeraati, et al., Attentional modulation of intrinsic timescales in visual cortex and spatial networks. *bioRxiv* (2021).
13. C Stringer, et al., Spontaneous behaviors drive multidimensional, brainwide activity. *Science* **364** (2019).
14. N Steinmetz, M Pachitariu, C Stringer, M Carandini, K Harris, Eight-probe Neuropixels recordings during spontaneous behaviors. (2019).
15. W Bair, E Zohary, WT Newsome, Correlated firing in macaque visual area mt: time scales and relationship to behavior. *Journal Neuroscience* **21**, 1676–1697 (2001).
16. D Levenstein, BO Watson, J Rinzel, G Buzsáki, Sleep regulation of the distribution of cortical firing rates. *Current Opinion Neurobiology* **44**, 34–42 (2017).
17. J De La Rocha, B Doiron, E Shea-Brown, K Josić, A Reyes, Correlation between neural spike trains increases with firing rate. *Nature* **448**, 802–806 (2007).
18. MD Humphries, JA Caballero, M Evans, S Maggi, A Singh, Spectral estimation for detecting low-dimensional structure in networks using arbitrary null models. *Plos one* **16**, e0254057 (2021).
19. P Sterling, S Laughlin, *Principles of neural design*. (MIT press), (2015).
20. JD Semedo, A Zandvakili, CK Machens, MY Byron, A Kohn, Cortical areas interact through a communication subspace. *Neuron* **102**, 249–259 (2019).
21. HA Swadlow, SG Waxman, Axonal conduction delays. *Scholarpedia* **7**, 1451 (2012) revision #125736.
22. JL Hennessy, DA Patterson, *Computer architecture: a quantitative approach*. (Elsevier), (2011).