Title: Core neurons are structural crossroads of cortical dynamics

Authors: Domenico Guarino[1]*, Anton Filipchuk[1,2], Alain Destexhe[1]*.

Affiliations:

[1] Department for Integrative and Computational Neuroscience (ICN), Paris-Saclay Institute of Neuroscience (NeuroPSI); Campus CEA Saclay, 151 route de la Rotonde, Bâtiment 151, 91400 Saclay, France.

[2] ICM, Institut du cerveau, Hôpital Pitié, 47 Bd de l'Hôpital, 75013 Paris.

*Corresponding author. Email: domenico.guarino@universite-paris-saclay.fr, alain.destexhe@cnrs.fr

Abstract: Cortical events of correlated neuronal firing are thought to underlie sensorimotor and associative functions. It is believed that events are attractors of cortical dynamics, pulled by strong mutual connections between recurrently active – core – neurons. Here, we exhaustively measured core structural connectivity, using a unique dataset with co-registered two-photon calcium imaging and electron microscopy of the same cortical tissue. We report that cores make fewer, weaker, mutual connections compared to other neurons. Instead, cores control the flow of the events they belong to. We present a model where cores arise from a simple distance-dependent connectivity rule, without imposed connectivity biases. Our integrated analysis shows that the identity of and connections between specific neurons, even when leading to attractor-like dynamics, cannot be neglected.

One-Sentence Summary: Recurrently active within reproducible population events, core neurons characterize cortical activity not by virtue of strong mutual connections but by controlling the flow of many crossing pathways.

Main Text:

According to Hebb (*1,2*), thought processes and behaviors correspond to the activation of *cell assemblies* – sets of neurons connected such that once a subset fires, the rest will fire. Their specificity and reliability would come from strengthening connections between coactive neurons. Inspired by the associative memory properties of assemblies, Hopfield (*3*) proposed a computational model, where strengthening connections between input-driven coactive units created *attractors* – reproducible activity patterns reliably associated with stimulus patterns. This model has become a cornerstone of cortical functioning (*4*), although convincing experimental evidence is available only for the hippocampus (*5*). In the past, correlation-based approaches have found traces of attractor dynamics in multi-unit recordings (*6*) but they were bounded by the limited number of recorded cells. Today, attractor-based models of cortical dynamics can be experimentally tested through the simultaneous recording of hundreds of cells, granted by optical physiology (*2,7*). Two-photon calcium imaging of the awake mouse cortex shows synchronous irregular population events, also called ensembles (*7*) or neural assemblies (*8*). These events are selectively evoked by stimuli, recur under the same stimulus presentations, keep their specificity when recurring spontaneously (*8–10*), and can be induced by repetitive coactivation (*11*). All these characteristics have reinforced the belief that population events are the result of local attractor networks (*7, 9–12*).

One feature of population events can be crucial to deciding whether they result from attractor networks. Recurring population events do not repeat identically. Their unifying feature is the *core* neurons – subsets of cells reliably participating in all event occurrences (*8, 9, 12*). Following the attractor paradigm, cores have been considered as biological equivalents of *pattern completion* units – capable of pulling the network dynamic towards a stored attractor, by virtue of strong mutual connections (*12, 19*). However, recent experiments show that repeated optogenetic co-activations do not strengthen synapses, as previously assumed (*11,12*), rather they increase cell excitability (*13*). Also, strong connections are theoretically associated with non-transient attractor states (*14*), slowing of network dynamics (*15*), and pathological run-away activity (*16*).

Here, we present evidence contrasting the hypothesis that core neurons are strongly connected. In particular, following the attractor paradigm, if local portions of the cortex show dynamic signs of storing attractors, like reproducible population events, they should also hold structural traces of attractors. We variously tested two fundamental attractor-driven assumptions: (i) *synapses between cores are more numerous and/or stronger compared to others*, and (ii) *circuits made by cores involve more recursive connections toward cores* (*3, 12, 17*). We found that core neurons do not make either more efficient or more recursive mutual connections. Instead, cores strategically sit at the confluence of connections crucial for the flowing of the events they belong to.

To perform the tests, we needed information about both the *firing* of neurons, to identify events and cores, and the *connections* between those neurons, to perform structural analysis of the circuits in which the cores were embedded. Until recently (*18*), this double requirement was impossible to fulfill. The MICrONS project (*19*) now gives this opportunity by releasing electron microscopy (EM) and co-registered two-photon calcium imaging datasets of the same cortical tissue. Two-photon recordings, and subsequent EM reconstruction, were made in layer 2/3 of the primary visual cortex of an awake mouse presented with static and drifting-oriented noise (*35*). The dataset has the duration and temporal resolution (1850 sec, at ~14.8 Hz), and structure (112 cells with both EM reconstruction and two-photon recording, with 1961 proofread synapses but

more than 3 million in total) to perform dynamical and structural analyses. Furthermore, we verified two-photon dataset generality against (N=32) recordings from the Allen Brain Observatory (*20*).

To identify recurrent spiking patterns, and their core neurons, we expanded a previously described method (*8, 9*) for the analysis of two-photon calcium imaging rasters (Fig. 1B). Briefly, a population event was identified as a peak of synchronous firing, beyond a surrogate-based significance threshold (Fig. 1C, and inset). The neurons firing within an event were listed in a vector (Fig. 1B, left). Pearson's correlation between event vectors was then used to gather events into clusters (Fig. 1D). Event pattern reproducibility – how much events were similar to one another within each cluster – was then established via another surrogate-based threshold (*35*). Finally, core neurons were singled out as those participating in the events of a cluster (Fig. 1D, inset) beyond a threshold justified below.

The MICrONS two-photon recording presented 226 significant events, and 10 clusters of significant reproducibility (>0.25), with 35 core neurons, reliably occurring across events (see example events in Fig. 1E). Events number, duration, and pattern reproducibility were similarly distributed in the MICrONS and Allen Brain Observatory datasets (Fig. S1).

Once cores were singled out, we proceeded with testing the attractor-driven assumptions. For the first – *synapses between cores are more numerous, and/or stronger, compared to others* – we took the number and volume of postsynaptic spines as a proxy for synaptic efficacy (*21-23*). Given the total of two-photon imaged cells and their proofread spines, the count of spines per core was unfortunately underpowered (*35*). Nonetheless, we used the current proofread dataset (constantly growing) to get an estimate of the probability of observing synapses, adopting several cautionary solutions: we lowered the threshold for core identification to a minimum (participation to 60% of the events), we did not differentiate spines based on their point of contact (e.g. axo-somatic, axo-dendritic, axo-axonic, etc), and we confirmed cores functional correlation (*24*, and Fig. S2). For each event, we grouped spines according to their core- or other-sources and targets. With the available frequencies, core neurons had a higher probability of having connections to and from non-cores than within their own group (Fig. 2A). The postsynaptic spine volumes made by core neurons were not bigger than those made by other neurons. There were also spines on cores and others (Fig. 2A, B, left), whose presynaptic neurons were either within or beyond the EM volume but they were not part of the two-photon dataset, therefore we did not integrate them into further analyses (but their statistics define the external input to our model network, below).

For the second assumption to test – *circuits made by cores involve more recursive connections towards cores* – we performed several graph theory analyses, considering cells as *nodes* and synapses as *edges* of the network (*25*). To get a global picture of mutual connectivity, we looked at the assortativity r – the proportion of mutual edges in a directed graph. We found r=-0.08, characteristic of non-mutually connected networks. To get a detailed picture, we looked at the distribution of motifs – statistically significant connectivity patterns considering groups of three cells (*26*). To understand the significance of motif occurrences, we used the ratio of real and 100 surrogate degree-matched network distributions (*27*). The ratio followed a known distribution (*27*) showing an abundance of three cell mutual motifs (Fig. 2C). However, looking at the motif distribution as a function of cores and other neurons across all population events of each cluster, we found that mutual, as well as convergent and divergent, motifs were not specific to cores (Fig. 2D).

Even if not having direct mutual connections, core mutual activity could be sustained by indirect synaptic feedback, via secondary or looped pathways, abundant enough to ensure that there were cliques of pattern completion units in traditional attractor network terms (*17*). These possibilities entail the existence of more edges, paths, and cycles involving cores than non-cores. We measured the number of cliques, length of shortest paths, and cycles between cores or non-cores. None gave a significant advantage to cores (Fig. S3).

The relatively large number of edges towards and from cores suggested that core neurons could be central to many cortical pathways. However, simple measures of centrality (*25*) gave no significant difference between cores and other neurons. The *degree* – number of edges incoming or outgoing a node – of cores was (non-significantly) higher than others (Fig. 2E). The *betweenness* – number of paths between two other nodes a neuron lies upon – of cores was not significantly higher than non-core neurons (Fig. 2F). The *hub centrality* – a measure of node centrality based on the centrality of the nodes it is connecting *to* – of cores was (non-significantly) higher than others (Fig. 2G).

The above measures were only concerned with the static – instantaneous – existence of edges between dynamically characterized nodes. However, population event dynamics is not instantaneous but rather a sequence of cell activations (Fig. 2H, left top, dashed red line). The cells in sequence might (or not) be directly connected (below, dashed red line). The *flow* algorithm (*28*), by weighting all possible edges between pairs of active cells in sequence, retrieves each cell's capacity to sustain that specific sequence. Cores had significantly more flow capacity than others (Fig. 2H, right). The counterpart of flow is the *cut* – a minimal set of edges that, if removed, would disconnect the sequence. Cores were on average significantly more targets of such edges (Fig. 2I). This result was confirmed by measuring the *PageRank* – which assigns scores to a neuron proportional to the scores of the neurons it is connected *by* (note the difference in direction with the hub score). Core neurons had significantly higher PageRank scores than others (Fig. 2J).

Altogether, there were no more numerous, or stronger, or more recursive connections between cores. But cores were central to the flow of population events' activity. Then, what is the connectivity principle leading to reproducible activity patterns and core neurons?

Based on the strong connectivity assumption, models have been proposed where strong recursive paths were arbitrarily imposed on randomly connected networks (*29*) by either additional circuits or increased synaptic strength between arbitrarily selected neurons (*12, 29, 30*). However, the predictive power of these models, concerning recurrent population events, is diminished by their ad-hoc solutions (*2*).

Here, we use a conductance-based balanced network model where the presence of core neurons results only from a distance-dependent connection strategy. In building our model, we established connections for each cell by randomly sampling its neighborhood with probability exponentially decreasing for increasing distances (Fig. 3A). We chose parameters such that connectivity loosely followed the EM distribution of connection distances and degrees (Fig. 3B, C). For model cell dynamics, we used already published parameters from mouse cortex (*31*). The model was driven by external independent Poisson spike generators (proportional in number to Fig. 2A non-local synapses).

The model produced synchronous irregular population events (Fig. 3D, E), onto which we applied the same core analysis as above (Fig. 3F), giving 431 significant events, of which 273 made 8 clusters of significant reproducibility, with 896 core cells.

In the model, we avoided entirely the first attractor-driven assumption, by assigning the same synaptic efficacy to all connections (and we tested results stability under several synaptic conductance values, Fig. S4). We then tested only the second attractor-driven assumption – core recursiveness – using the above graph measures analysis.

The distance-dependent connection strategy was sufficient to get core vs other distributions for degree, betweenness, and flow, resembling in significativity and effect size the relationships found in the MICrONS dataset (Fig. 3G-L). As in that dataset, the model presented no evidence of more mutual connectivity via indirect synaptic feedback for cores (Fig. S5).

The distance-dependent connection strategy was also necessary. When the same network was randomly rewired, preserving node degrees, it gave rise to population events, but no reproducible pattern was identified, and consequently no cores (Fig. S6).

Without imposing strong recursive connections (*12, 29*), our modeling confirmed that locally clustered (here via distance-dependent) connectivity produces weakly recursive circuits that irregularly transition between periods of higher and lower firing rates (*30, 32, 33*). Importantly, we add that distance-dependent connection clustering is also the source of population event reproducibility and cores.

There are several limitations in this study: (i) a single mouse dataset necessarily gives a reduced and biased picture; (ii) its two-photon/EM sample size is still limited; and (iii) sensory cortical areas are less likely to express attractor-like dynamics compared to prefrontal areas (*17, 34*). To overcome these limitations, it can be mentioned that the coordinated EM/two-photon approach will generate more datasets, improving the present picture. In addition, we supply a model built using very conservative assumptions, which closely mimicks structural and dynamical experimental features. Besides, the general distance-dependent strategy we adopted is also susceptible to varying sampling distributions beyond exponential (*33*), which can be adjusted to represent different cortical areas.

Uniformity and modularity have been used to argue for an attractor-like nature of the cortex (*2, 4, 12*). Accordingly, the role of cortical networks in cognition would be better understood as transformations between attractor representations, where the identity of and connections between specific neurons can be safely regarded as details (*4*). However, when we looked at local portions of the cortex showing dynamic signs of attractors – like reproducible population events and cores – and searched for structural traces of attractors – like enhanced connectedness of core neurons – we found none. Instead, the ability of specific neurons to sustain the flow of activity along event sequences cannot be regarded as detail and vouches for tight integration between the specificity of local circuits and the attractor-like properties of larger portions of the cortex.

References and Notes

1. D. O. Hebb, *The organization of behavior: A neuropsychological theory* (Psychology Press, 2005).
2. R. Yuste, From the neuron doctrine to neural networks. *Nature reviews neuroscience,* ***16*** (8), 487-497 (2015).
3. J. J. Hopfield, Neural networks and physical systems with emergent collective computational abilities. *Proc. Natl. Acad. Sci. U.S.A.,* ***79***, 2554-2558 (1982).

4. D. L. Barack, J. W. Krakauer, Two views on the cognitive brain. *Nature Reviews Neuroscience, **22***, 359-371 (2021).

5. T. Wills *et al.*, Attractor Dynamics in the Hippocampal Representation of the Local Environment. *Science, **308*** (5723), 873–876 (2005).

6. R. Cossart, Y. Ikegaya, R. Yuste, Calcium imaging of cortical networks dynamics. *Cell calcium, **37***, 451-457 (2005).

7. Y. Prut, *et al*. Spatiotemporal structure of cortical activity: properties and behavioral relevance. *Journal of neurophysiology*, ***79***(6), 2857-2874, (1998).

8. A. Filipchuk, A. Destexhe, B. Bathellier, Awake perception is associated with dedicated neuronal assemblies in cerebral cortex. https://www.biorxiv.org/content/10.1101/2021.08.31.458322v1 (2021).

9. J. Miller, *et al*., Visual stimuli recruit intrinsically generated cortical ensembles. *Proc. Natl. Acad. Sci. U.S.A.*, **111**, E4053-E4061 (2014).

10. L. Carrillo-Reid, *et al*., Controlling visually guided behavior by holographic recalling of cortical ensembles. *Cell, **178***, 447-457 (2019).

11. L. Carrillo-Reid, *et al*., Imprinting and recalling cortical ensembles. *Science, **353***, 691-694 (2016).

12. L. Carrillo-Reid, *et al*., Identification of Pattern Completion Neurons in Neuronal Ensembles Using Probabilistic Graphical Models. *Journal of Neuroscience, **41***, 8577-8588 (2021).

13. A. Tzitzitlini, *et al.*, Intrinsic excitability mechanisms of neuronal ensemble formation. https://www.biorxiv.org/content/10.1101/2020.07.29.223966v2 (2020).

14. Y. Wu, F. Zenke, Nonlinear transient amplification in recurrent neural networks with short-term plasticity. *Elife, **10***, e71263 (2021).

15. S. Sadeh, C. Clopath, Inhibitory stabilization and cortical computation. *Nature Reviews Neuroscience, **22***, 21-37 (2021).

16. D. Zhang, *et al.*, Global enhancement of cortical excitability following coactivation of large neuronal populations. *Proc. Natl. Acad. Sci. U.S.A. **117***, 20254-20264 (2020).

17. E. T. Rolls, Attractor networks. WIREs Cognitive Science, 1, 119–134 (2010).

18. D. Bock, *et al*., Network anatomy and in vivo physiology of visual cortical neurons. *Nature, **471***, 177-182 (2011).

19. S. Dorkenwald et al. Binary and analog variation of synapses between cortical pyramidal neurons. https://doi.org/10.1101/2019.12.29.89031 (2019).

20. S. de Vries, *et al*., A large-scale standardized physiological survey reveals functional organization of the mouse visual cortex. *Nature Neuroscience **23***, 138–151 (2020).

21. A. Vanharreveld, E. Fifkova, Mechanism for potentiation and short-term memory. *Proceedings of the Koninklijke Nederlandse Akademie, **78***, 21-24 (1975).

22. M. Matsuzaki, *et al*., Dendritic spine geometry is critical for AMPA receptor expression in hippocampal CA1 pyramidal neurons. *Nature neuroscience, **4***, 1086-1092 (2001).

23. S. Barnes, *et al.*. Deprivation-induced homeostatic spine scaling in vivo is localized to dendritic branches that have undergone recent spine loss. *Neuron,* ***96***, 871-882 (2017).

24. A. Sadovsky, J. MacLean, Scaling of topologically similar functional modules defines mouse primary auditory and somatosensory microcircuitry. *Journal of Neuroscience,* ***33***, 14048-14060 (2013).

25. M. Newman, *Networks – An Introduction* (Oxford University Press, 2010).

26. R. Milo, *et al.*, Network motifs: simple building blocks of complex networks. *Science,* ***298***, 824-827 (2002).

27. S. Song, *et al.*, Highly nonrandom features of synaptic connectivity in local cortical circuits. *PLoS Biology,* ***3***, e68 (2005).

28. A.V. Goldberg et al., Network flow algorithms, *Tech. Report STAN-CS-89-1252*, (1989).

29. T. Vogels, L. Abbott, Signal propagation and logic gating in networks of integrate-and-fire neurons. *Journal of Neuroscience,* ***25***, 10786-10795 (2005).

30. A. Litwin-Kumar, B. Doiron, Slow dynamics and high variability in balanced cortical networks with clustered connections. *Nature Neuroscience,* ***15***, 1498-1505 (2012).

31. Y. Zerlaut, *et al.*, Heterogeneous firing rate response of mouse layer V pyramidal neurons in the fluctuation driven regime. *The Journal of physiology,* ***594***, 3791-3808 (2016).

32. P. Yger, *et al.*, Topologically invariant macroscopic statistics in balanced networks of conductance-based integrate-and-fire neurons. *Journal of computational neuroscience*, ***31***, 229-245 (2011).

33. S. Spreizer, A. Aertsen, A. Kumar, From space to time: Spatial inhomogeneities lead to the emergence of spatiotemporal sequences in spiking neuronal networks. *PLoS computational biology,* ***15***(10), e1007432 (2019).

34. H. K. Inagaki, Discrete attractor dynamics underlies persistent activity in the frontal cortex. *Nature* ***566***, 212–217 (2019).

35. Materials and methods are available as supplementary materials at the Science website.

Methods references

36. G. Csárdi, T. Nepusz, The igraph software package for complex network research. Inter *Journal Complex Systems,* ***1695*** (2006).

37. D. Kätzel, et al., The columnar and laminar organization of inhibitory connections to neocortical excitatory cells. *Nature Neuroscience,* ***14***, 100-107 (2011).

38. H. Markram, et al. Reconstruction and simulation of neocortical microcircuitry. *Cell,* ***163***, 456-492 (2015).

39. S. C. Seeman, et al., Sparse recurrent excitatory connectivity in the microcircuit of the adult mouse and human cortex. *Elife,* ***7***, e37349 (2018).

40. M. Diesmann, M. Gewaltig, NEST: An environment for neural systems simulations. *Forschung und wisschenschaftliches Rechnen, Beiträge zum Heinz-Billing-Preis,* ***58***, 43-70 (2001).

41. A. Davison et al., PyNN: a common interface for neuronal network simulators. *Frontiers in neuroinformatics, **2***, 11, (2009).

42. R. Wilcox, H. Keselman, Modern robust data analysis methods: measures of central tendency. *Psychological Methods*, **8**(3), 254–274, (2003).

Acknowledgments: We express our gratitude to the MICrONS project and the Allen Brain Institute for the efforts they make and their commitment to publicly releasing their datasets, without which this study would have not been possible. We would like to thank Yves Frègnac, Manuel Mameli, Federico Tesler, and Jan Antolik, for stimulating discussions and critically reading the manuscript. The mouse drawing in Fig. 1 is kindly offered by Ethan Tyler and Lex Kravitz.

Funding:

EC Human Brain Project (HBP, grant agreement no. H2020-945539).

Author contributions:

Conceptualization: DG, AF, AD

Methodology: DG, AF

Investigation: DG, AF, AD

Visualization: DG, AF

Funding acquisition: AD

Project administration: AD

Supervision: AD

Writing – original draft: DG, AF, AD

Writing – review & editing: DG, AF, AD

Competing interests: Authors declare that they have no competing interests.

Data and materials availability: All data, code, and materials used in the analysis are publicly available to any researcher for purposes of reproducing or extending the analysis. The MICrONS project dataset is available at: https://www.microns-explorer.org/phase1. The Allen Brain Observatory dataset is available at: http://observatory.brain-map.org. The calls to retrieve the Allen data are listed in the Materials and Methods section. The code to run all analyses is available at https://github.com/dguarino/Guarino-Filipchuk-Destexhe. An interactive Jupyter notebook is accessible on Binder: https://mybinder.org/v2/gh/dguarino/Guarino-Filipchuk-Destexhe/HEAD. The code to run all simulations (complete with Dockerfile to recreate the environment) is available at https://github.com/dguarino/Guarino-Filipchuk-Destexhe-simulations.

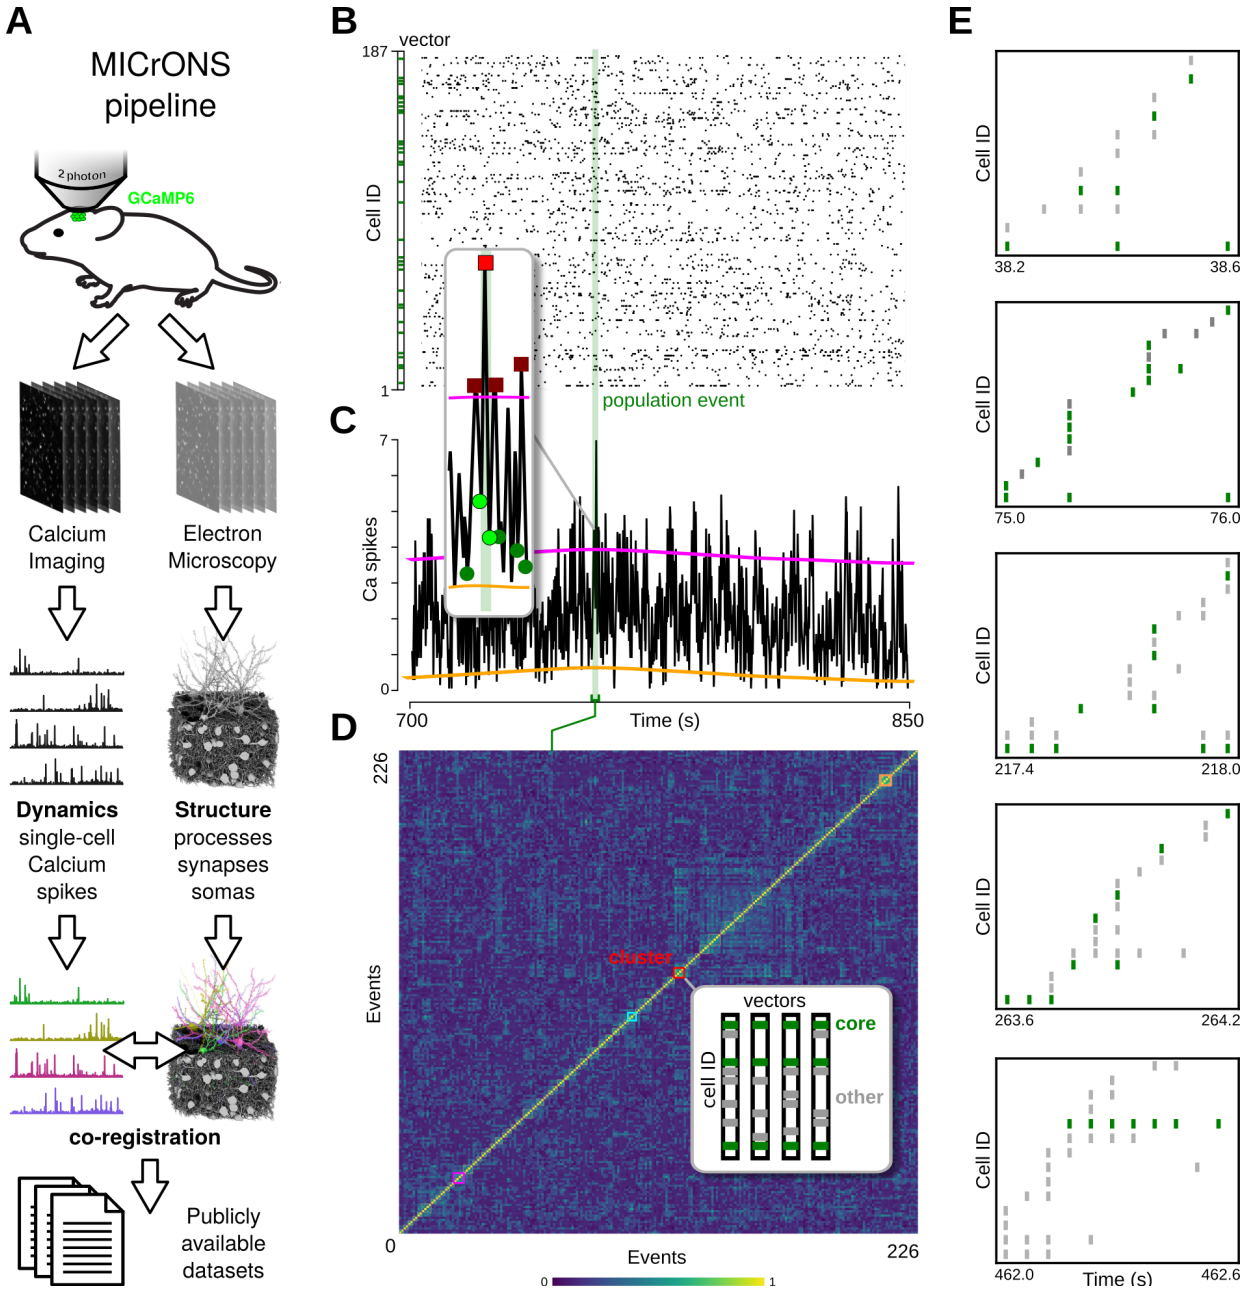

Fig. 1. Population events and core neurons. (**A**) Schematic diagram of the MICrONS pipeline. In-vivo calcium imaging and serial EM imaging were performed on the same mouse cortical volume, yielding electrophysiological dynamics and spatial structure of neurons and synapses. An automatized co-registering procedure matched each neuron visible in the calcium imaging movie with a corresponding neuron in the anatomical EM stack. After human proofreading, the data was released to the public. (**B**) Calcium spikes raster plot from awake mice of the MICrONS dataset. Fluctuations in spike frequency indicate synchronous population events (example green shade). (**C**) Population events were quantified as time intervals where the firing rate (black) exceeded a surrogate-based threshold (magenta) above firing baseline (orange). *Inset*. The subthreshold beginning and end of each event were the minima (green points) before and after each peak (red squares). Each event was assigned a population binary vector (**B**, participating

cells in green). (**D**) Each population binary vector was correlated with all others (see color bar for degree of correlation) and clustered based on their similarity. Population events sharing neurons beyond a surrogate-based threshold were identified (colored squares), their autocorrelation computed, and the cells reliably participating in their events were labeled as "cores" (Right, inset). (**E**) Example events from the cluster highlighted in **D**. Cells are displayed sorted by time of occurrence.

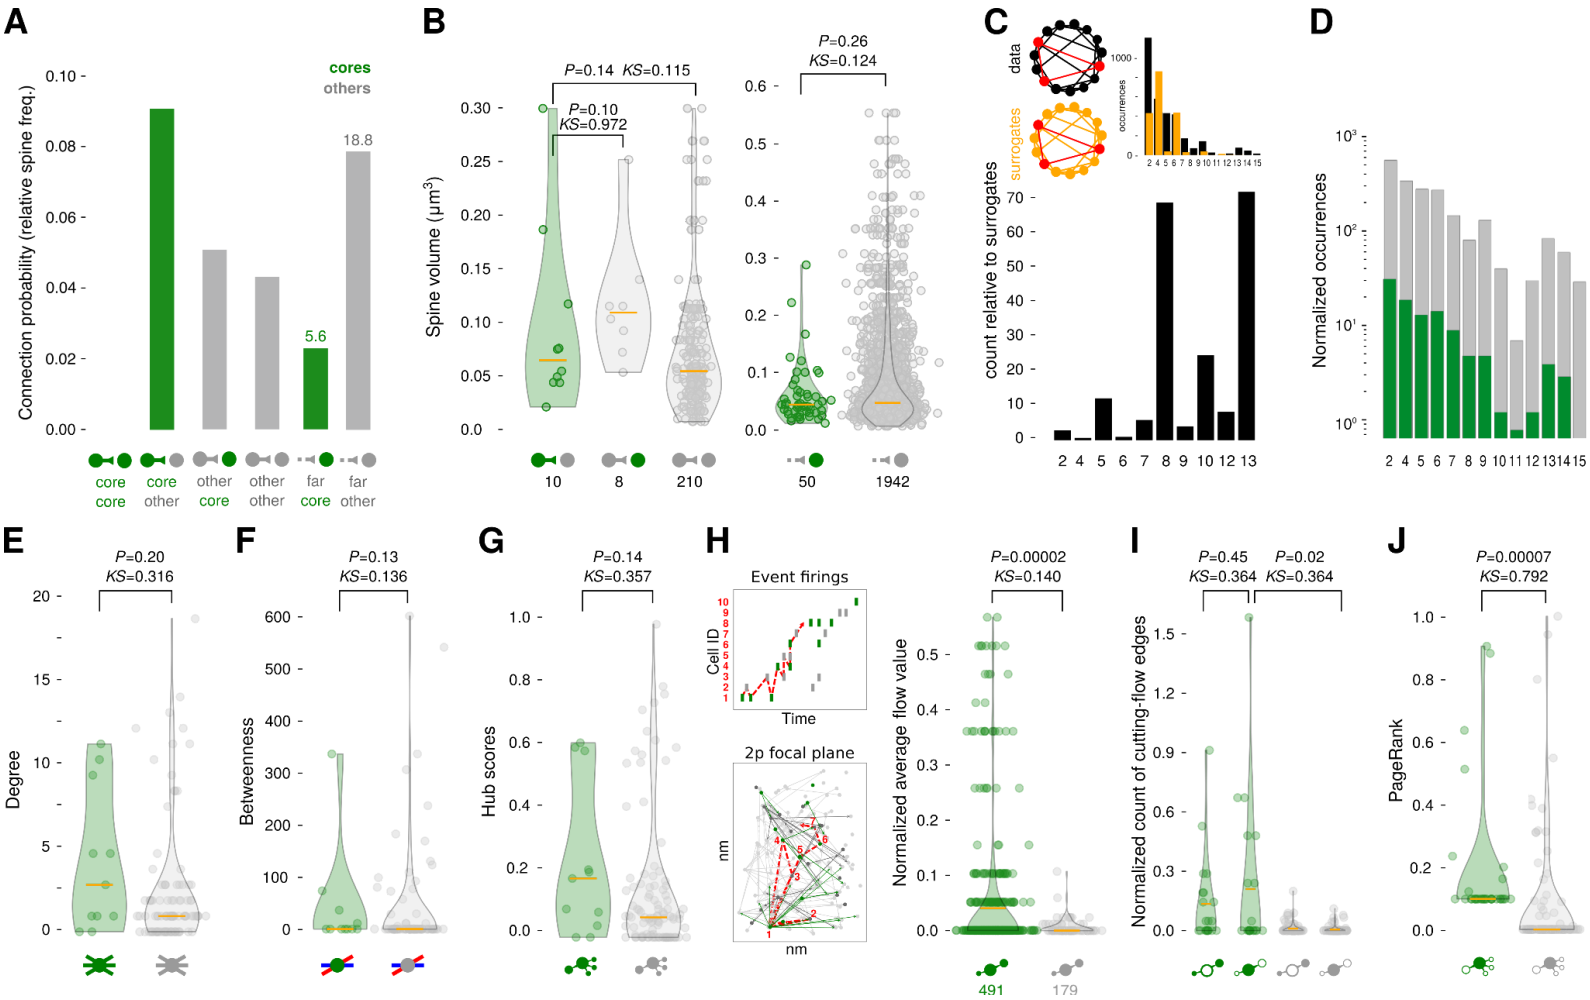

Fig. 2. Core neurons are crossroads of cortical structure and dynamics. (**A**) Connection probability for core (green) and other non-core (gray) neurons. Note the absence of core-to-core synapses. (*Right*) Normalized count of synapses from somas beyond EM volume. (**B**) The distributions of postsynaptic spine volumes between cores and others were not significantly different. (*Right*) Postsynaptic spine volumes from neurons beyond the scanned tissue were not innervating significantly more cores or others. (**C**) EM connectivity motifs distribution for all two-photon recorded neurons (*top*: all motifs, numbered for clarity, exceeded surrogate networks). (**D**) Normalized motif occurrence distribution (numbers stand for standard motifs) for cores and other neurons. Both made a mixture of mutual and non-mutual connections. (**E**) Core neurons had a (non-significant) larger number of connections compared to other non-cores. (**F**) Core neurons were not on more pathways than others. (**G**) Cores were (non-significantly) more connected to high-degree nodes. (**H**) Event flow analysis. (*Left*) The sequence of activations (*top*) does not necessarily correspond to its cells subnetwork (*below*). By following all possible pathways between active cells in sequence, we retrieve each cell's flow capacity. (*Right*) Cores had significantly more (normalized) flow capacity. (**I**) Removing connections having cores as targets (green, *right*) stopped significantly more flows than targeting non-cores. (**J**) Core neurons had significantly higher PageRank scores than non-cores.

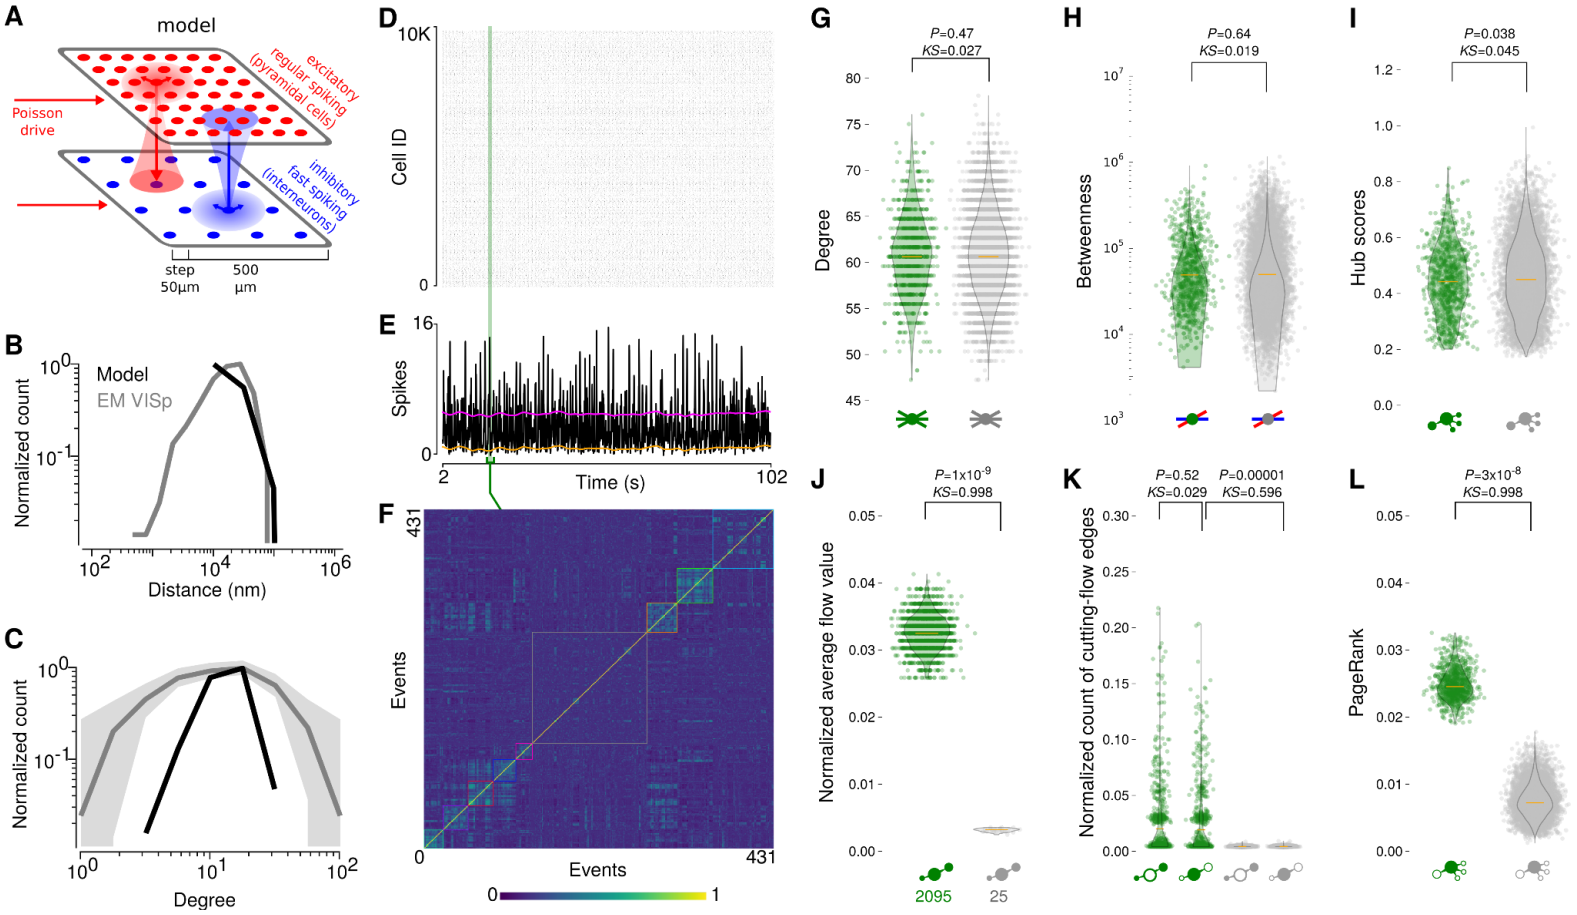

Fig. 3. A spiking model reproduces in-vivo core features. (**A**) The model was made of 10000 excitatory and 2500 inhibitory adaptive-exponential integrate-and-fire point neurons on a grid of approximately 500µm with units separated by 5µm. (**B**) The spatial reach of model distance-dependent connections matched the EM distribution of distances. Due to the hard-coded separation of model neurons, the smallest connection diameter in the model was 10µm. (**C**) The model connectivity rule induced a degree distribution within the interval of the EM data. (**D**) Spikes raster plot with synchronous irregular population events (example green shade). The model showed a dynamic similar to that observed in the 2p dataset. (**E**) Population events were quantified as in fig. 1A. (**F**) Population binary vectors were correlated with all others and clustered as in fig. 1C (gray squares identify non-significant clusters). (**G**) The distribution of cores number of connections was not different from other neurons. (**H**) Cores were not on more pathways than others. (**I**) Cores were not preferentially connected to highly connected neurons. (**J**) Core neurons sustained significantly more than others the flow of the event sequences they participated in. (**K**) Removing core connections from the network graph stopped more flows from beginning to end of event sequences. (**L**) Core neurons had significantly higher PageRank scores than non-cores.

Supplementary Materials for

Core neurons are structural crossroads of cortical dynamics

Domenico Guarino, Anton Filipchuk, Alain Destexhe

Correspondence to: domenico.guarino@universite-paris-saclay.fr, alain.destexhe@cnrs.fr

This PDF includes:

Materials and Methods

Fig. S1 to S6

Materials and Methods

We performed the same structural and dynamical analyses over data collected in mouse cortex and in biologically constrained conductance-based spiking models. All animal procedures were in accordance with the Institutional Animal Care and Use Committees at the Baylor College of Medicine and the Allen Institute for Brain Science.

MICrONS data source. The MICrONS project (https://www.microns-explorer.org/) made available, through its website, two datasets, corresponding to the first two phases of the project. The first phase, the one we used, has electron microscopy and functional imaging of the same cortical patch of V1, in a 400 400 200 µm volume with the superficial surface of the volume at the border of L1 and L2/3, approximately 100µm below the pia. Here follows a brief description of the procedures, and we refer the interested reader to Dorkenwald et al. 2019. **Two-Photon imaging**. Functional imaging was performed in a transgenic mouse expressing fluorescent GCaMP6f from the Jackson Laboratories (jax.org) and the Allen Brain Institute (see next section on the Allen data source). The complete description of mouse transgenic line, surgical procedures, implants, two-photon microscope features, stimulation and recording setups, and protocols were already published (Dorkenwald et al. 2019). One important feature and the main difference with the Allen Brain Observatory two-photon dataset is the recording acquisition frame rate, here at ~14.8 Hz, resulting in time-frame windows of ~67 ms (the Allen dataset has double this resolution, ~30 Hz). **Stimuli**. The mouse was head-restrained but could walk on a treadmill during imaging. Visual stimuli were 30 one-minute trials of a colored-noise interspersed with periods of coherent motion of oriented noise. Each one-minute trial contained 16 stationary-moving-stationary blocks, with a different direction presented in each block, pseudorandomly-ordered. **Electron microscopy**. After functional imaging, the mouse was perfused with fixant agents, dissected, and 200 µm coronal slices of the V1 area, interested by two-photon sessions, were cut. Serial sections of 40 nm thickness were cut from the slices and prepared for transmission electron microscopy. Custom modifications to the setup and camera allowed fields-of-view as large as (13 μm^2) at 3.58 nm resolution. Data were acquired in blocks, aligned through multiple steps, and defects manually detected. Somas, synaptic cleft, and spine detection and partnering were manually performed on a (387 clefts) subset for training and testing machine learning libraries. The rest were performed by machine learning libraries.

Finally, cells identified in the two-photon imaging max projection were manually co-registered with EM-identified somas.

Allen data source. The Allen Institute made available, through its Brain Observatory, a large dataset of neuron responses from several visual cortical areas and layers, using mouse transgenic lines expressing a Cre/Tet-dependent fluorescent calcium indicator (GCaMP6f), which allows the recording of calcium influx associated with neural activity. The complete description of mouse transgenic lines, surgical procedures, behavior training, implants, two-photon microscope features, stimulation and recording setups, and protocols were already published (see Allen VisualCoding_Overview whitepaper in the references). Here we briefly provide the relevant information in the context of our work. **Cell type and experiment selection**. The Allen Brain Observatory contains data collected in several cell types of mouse visual cortices (see Allen VisualCoding_Genotyping whitepaper in the references). Cell-type specific expression of GCaMP6f was conferred by the Cre recombinase and tetracycline-controlled transactivator protein (tTA) under the control of the Camk2a promoter with the Ai93 and Ai94 reporters. To maximize the number of recorded neurons across layers, we selected all Cre drivers specific for pyramidal neurons and layers. Considering that each layer of mouse V1 is roughly 100um thick (Scala et al. 2019), we were able to design cell-type and layer-specific queries which resulted in the selection of the following experiment ids (as of 2019/09/10). We selected layer 2/3 planes, in order to have a dataset comparable to the one from the MICrONS project (query: http://observatory.brain-map.org/visualcoding/search/overview?area=VISp&imaging_depth=175,185,195,200,205,225,250,265,275,276,285,300&tld1_name=Slc17a7-IRES2-Cre,Cux2-CreERT2,Emx1-IRES-Cre), selected experiment ids: [502115959, 501574836, 502205092, 502608215, 510514474, 503109347, 510517131, 524691284, 645413759, 660513003, 704298735, 702934964, 712178511, 526504941, 528402271, 540684467, 545446482, 561312435, 596584192, 647155122, 652094901, 652842572, 653122667, 653125130, 657080632, 661328410, 661437140, 663485329, 679702884, 680156911, 683257169, 688678766] (with exp: 501271265 removed due to https://github.com/AllenInstitute/AllenSDK/issues/66). A total of 32 experiments were selected and downloaded as NWB files. Each experiment contained pre-processed calcium indicator fluorescence movies (512x512 pixels, i.e. 400 µm cortical field of view, sampled at 30Hz), corrected for pixel leakage during scanning, ROI filtered, and

neuropil subtracted. Each resulting NWB file consisted of an experiment session (~90 min) at a given location (cortical area and depth).

Calcium spike timing extraction. To infer the action potential timings from the calcium indicator fluorescence of the Allen Brain data source, we used the python library OASIS (Friedrich et al. 2017). This library for spike inference, which is based on an active-set method to infer spikes, also has routines for automatic estimation of its parameters. The library is available at https://github.com/j-friedrich/OASIS. As parameters for the deconvolve function, we used *penalty*=1, and we used the automatic estimation of fluorescence variation from the data (parameter *g*). The resulting spike timings obtained with these parameters were verified for adherence to the example deconvoluted data provided by the Allen Brain Observatory ipython notebooks.

Dynamical analysis. For each data source, we applied the same dynamical analysis. The analysis of synchronous irregular population events has already been developed in MATLAB™ and described in Filipchuk et al. 2022 (itself derived from Miller et al. 2014, and before going back to Prut et al. 1998), please refer to it for an in-depth description. We reimplemented all analysis steps in Python and tested them against the original analysis. Briefly, single cell activity was described as a cell binary vector, where the number of elements was equal to the T, number of time frames during the recording (frame duration was ~30ms). A value of 1 was assigned at the time frames corresponding to the onset of each spike, and the cell vector was zero otherwise. The number of cell vectors was equal to the total number of recorded neurons in the field of view, N_e, resulting in a TxN_e spiketrain matrix where columns corresponded to the time frames and rows to the neurons. Each column was summed to give the number of neurons coactive at every time frame from which the instantaneous population firing rate could be deduced. As is visible from sample raster plots (Fig. 1A), there were periods of time where spikes were much more synchronized across the population, which could be interpreted as population events, departing from the fluctuations of an asynchronous population spiking process. To identify whether the population events were above chance coincidence, and determine their exact start and end, we applied the following algorithm. (**1**) We created 100 surrogate TxN_e spiketrain matrices by reshuffling each neuron's inter-spike interval. The average 99 percentile across time frames was then calculated and a firing rate threshold for event significance was obtained by adding this average value to a local baseline estimate that aimed at correcting slow fluctuations

of background firing rate. The baseline estimate was generated using an asymmetric least squares smoothing algorithm from Eilers and Boelens (2005), with parameters $p=0.01$ for asymmetry and $\lambda=10^8$ for smoothness. (**2**) Experimental population firing rate trace was smoothed using Savitzky-Golay filtering, with order=3 and frame window=7, to get rid of non-essential peaks. Local maxima for the smoothed curve that were above the previously defined threshold were retained (Fig. 1B, inset, red square). The adjacent local minima around each of the local maxima were identified (Fig. 1B, inset, green circles). Their time frames were taken as the beginnings and ends of population events (Fig. 1A, limits of the green vertical band). (**3**) The time interval between the beginning and end of each population event defined its duration. We described each population event as a binary vector of length equal to the total number of neurons in the recorded field of view (N_e). We initialized all vector's elements to *0*, and then we indicated the participation of a cell to the event, spiking at least once over its duration, with a coordinated *1* (Fig. 1A, left). (**4**) Pearson's correlation matrix for ongoing and evoked events was constructed by computing the Pearson correlation coefficient between all binary vectors. In plotting the matrix, the color reflected the degree of correlation with a colormap ranging from 0 to 1. (**5**) Population events sharing a substantial number of neurons, i.e. having a significant degree of correlation, were clustered together (using agglomerative linkage clustering with the farthest distance based on correlation metric, from the python *scipy.clustering* package), and the correlation matrix could be permuted according to the clustering order (Fig. 1C). (**6**) The correlation between population events might be simply due to the finite number of recorded cells. We, therefore, defined a significance threshold for intra-cluster correlation, established using surrogate events. We generated $100 \times N_e$ surrogate event vectors (with the same number of cells per event of the originals, but randomly chosen cells, without replacement). We then performed correlation-based clustering (as in point 5) of the resulting surrogate clusters, and we used the 95 percentile of correlation coefficient of these clusters as a correlation threshold. (**7**) To quantify the pattern reproducibility of events within a cluster, i.e. the number of cells shared across events in a cluster, the correlation was calculated across all events of each cluster. (**8**) Core cells were identified as those participating in a range of thresholds (from 55%, barely above no difference in participation, to 95%) of the events within each cluster. Relaxing this requirement to allow for a more participative definition of cores changed the statistics but not the trends presented in Fig. 2. Once identified, we verified that the average firing rate of each core was higher within than that outside its events, discarding those with unspecific firing. In developing

our dynamical analysis for the MICrONS data, we used the notebook example on the Allen Institute (https://github.com/AllenInstitute/MicronsBinder/blob/master/notebooks/vignette_analysis/function/structure_function_analysis.ipynb). The whole analysis workflow code, is available on GitHub (https://github.com/dguarino/Guarino-Filipchuk-Destexhe), including also an interactive Jupyter Notebook version on Binder for faster reproducibility (https://mybinder.org/v2/gh/dguarino/Guarino-Filipchuk-Destexhe/HEAD).

Structural analysis. We used the data tables published by the MICrONS project on their website and the referenced publications (see above). **Postsynaptic spine number and volume analysis**. We selected the subset of proofread synapses having as soma root id those identified and co-registered in the two-photon dataset. After the dynamical analysis was performed (see above), core neuron ids for each event were available and used to query the synapse table, and postsynaptic spine volumes were collected by identified type (core or non-core). The four possible combinations (core-core, core-other, other-core, other-other) were plotted in Fig. 2. Given that different cores participated in different clusters, the number of cores and others differed, and each core (or non-core) neuron could potentially be connected to all other cores (or non-cores) we reported in Fig. 2A the normalized number of spines. Differences in significance and effect size were established using Kruskal-Wallis and Kolmogorov-Smirnov distance, since the distributions were non-normal, as reported below in the statistical methods. **Synaptic functional efficacy analysis**. To estimate functional connectivity, we adopted the same method described in Sadovsky and MacLean (2013). Briefly, an adjacency matrix representing a graph for each experiment was created with nodes being all recorded cells, and directed edges formed according to a single frame lagged correlation (frame duration: 67.4ms) and weighted according to the number of spikes (see supplementary figure 1). **Graph theory analyses**. An adjacency matrix was built by coordinating pre- and post- soma root ids. All graph analyses reported in Fig. 2 were performed using the python version of the library *iGraph* (Csárdi and Nepusz 2006). The iGraph library allows motif identification with either 3 or 4 nodes. We used 3-motifs to be able to compare our results with other publications and with the online MICrONS ipython notebooks, and also to make computations of the 100 surrogate graphs run in a reasonable time. The comparison of core-based cycles and other-based cycles was performed over an adjacency list and using a function that recursively checked over the list of neighbors starting with each core

(or non-core) cell id. In developing our structural analysis for the MICrONS data, we used the notebook example on the Allen Institute GitHub page (https://github.com/AllenInstitute/MicronsBinder/blob/master/notebooks/intro/MostSynapsesInAndOut.ipynb). The whole analysis workflow code, including also an interactive Jupyter Notebook version on Binder, is available (see above).

Statistical methods. The data released during phase 1 by the MICrONS project concerned one mouse. This limit was our major concern while designing our study. However, our unit of statistical analysis was the number of events. Events are produced by the same cortical tissue overall but by different cells and connections. This implies that we did not set out to measure the same outcome multiple times on the same animal – which would entail an assay variability. Instead, we set out to measure different (quasi-)independent outcomes – which is inherent biological variability (the high level of interconnectivity does not make it completely independent). This is why, in the currently released MICrONS dataset, we consider as statistically legitimate our sample size of n=226 events, i.e. the number of independent observations of our unit of analysis, under a single experimental condition. The goal of our study was hypothesis testing, therefore we designed each analysis as a statistical hypothesis test to produce an exact statistical significance level (*P-value*) with an appropriate effect size. Concerning the attractor features for cores vs non-cores, they are taken to exist if the following hold in the specific direction (one-tailed direction): (1) Core neurons have more numerous synapses between themselves, or towards others, compared to those of non-cores. (2) Spines with pre-synaptic cores and post-synaptic cores or non-cores are bigger than spines with pre-synaptic non-cores. The required sample size, common to both points, can be computed by considering the following. The study group ("cores") is compared to the rest of the population (non-cores, or "others"). The primary endpoint is a continuous value (synaptic spine volume, μm^3). The anticipated mean of the "others" is 0.02±0.5 (based on the recorded average). Estimating a higher average for cores of 0.08±0.5, with a type I error probability (false positive rate) of α=0.05, a type II error probability (false negative rate) of β=0.3 (corresponding to a modest power of 70%), the resulting sample size for cores is N=254. The actual number of two-photon recorded cells is N=112. Therefore, in the current status of the MICrONS project, the spine count for cores is underpowered. In order to provide an estimate, we considered two aspects: (a) given that the spine count for each cluster of events with its cores and others is taken independently,

many spines will be considered in a repeated measure fashion; (b) the maximum value of the independent variable (spine volume) is >10 times its minimum value (respectively 0.35 and 0.02 μm^3), allowing a large space of possible values. For all tests, to evaluate the differences between distributions resulting from our measurements of the nominal variables (spine volumes, 1-lag correlations, degrees, betweenness, …) we considered that: the nature of the outcome was continuous, the number of groups was larger than 2 (unless specified, or with unequal measurements), the groups were independent (cores and non-cores groups do not share cells). We tested distributions' normality using the D'Agostino and Pearson's test. For normally distributed variables, we used the One-way ANOVA test. For non-normally distributed variables, we adopted the Kruskal-Wallis test. For some variables, the sample size was large (beyond 100 samples), therefore even small variations could be labeled as significant by the above tests. In these cases, a quantitative measure of the effect size was required. As a measure of effect size, we adopted the Kolmogorov-Smirnov test difference, which is more robust than Cohen's delta for non-normal distributions (Wilcox and Keselman 2003). For all these tests, we used the Python *scipy* library.

Neuron model. All simulated networks were composed of n=12500 (10000 excitatory and 2500 inhibitory) conductance-based adaptive exponential integrate-and-fire (AdExIF) point neurons. The set of equations describing the membrane potential (1) and spike adaptation current (2) of a neuron are

$$C_m \frac{dV}{dt} = -g_L(V - E_L) + g_L \Delta \, exp[(V - V_T)/\Delta] - w/S \quad (1)$$

$$\frac{dw}{dt} = \frac{1}{\tau_w}[a(V - E_L) - w] \quad (2)$$

with C_m being the membrane capacitance (nF), g_L the resting (leak) conductance (μS), E_L the resting (leak) potential (mV, which is also equal to the reset value after spike), Δ the steepness of the exponential approach to the threshold (mV), V_T the spike threshold, and S membrane area (μm^2). When V reaches the threshold, a spike is emitted and V is instantaneously reset and clamped to the reset value during a refractory period (ms). w is the adaptation variable, with time constant τ_w (ms), and the dynamics of adaptation is given by parameter a (in μS). At each spike, w is incremented by a value b (in nA), which regulates the strength of adaptation. An

experimentally informed parameter set was chosen (from Zerlaut et al. 2011) and kept fixed, see Table 1.

Parameter	Value	
Regular Spiking (excitatory)		
Threshold (V_T)	-50.0	mV
Resting potential (E_L)	-65.0	mV
Afterpotential reset	-65.0	mV
Refractory time constant	5	ms
Membrane time constant (τ_m)	15.0	ms
Membrane capacitance (C_m)	0.15	nF
Subthreshold adaptation (a)	4.0	nS
Spike-triggered adaptation (b)	0.02	nA
Slope factor (Δ_T)	2.0	mV
Adaptation time constant (τ_w)	500	ms
Fast Spiking (inhibitory)		
Threshold (V_T)	-50.0	mV
Resting potential (E_L)	-65.0	mV
Afterpotential reset	-65.0	mV
Refractory time constant	5	ms
Membrane time constant (τ_m)	15.0	ms
Membrane capacitance (C_m)	0.15	nF
Subthreshold adaptation (a)	0.0	nS
Spike-triggered adaptation (b)	0.0	nA
Slope factor (Δ_T)	0.5	mV
Adaptation time constant (τ_w)	500	ms

Table 1. Parameters of the AdEx Integrate-and-Fire model.

Synapse model. The synaptic connections between neurons were modeled as transient conductance changes. The synaptic time course was modeled as an alpha function with a fast rise followed by an exponential decay:

$$\tau_{syn}\frac{dg_{syn}(t)}{dt} = -g_{max}(t - S_{syn}(t))\, exp[-(t - S_{syn}(t)) - \tau_{syn}] \quad (3)$$

where $syn \in \{exc, inh\}$, $S_{syn}(t)=\delta(t - t_k)$ are the k incoming synaptic spike trains where $i \in \{1, .., N\}$ refers to presynaptic neurons and k to the different spike times of these neurons. The synaptic time constants were chosen to be τ_{exc}=3 ms and τ_{inh}=7 ms for excitation and inhibition respectively. The reversal potentials were E_{exc}=0 mV and E_{inh}=−80 mV. The synaptic weights

$g_{syn}(t)$ were set to g_{exc}=1 nS for the excitatory conductance, and a balance of g_{inh}=5 nS unless stated otherwise. We used a constant delay of 0.5 ms, equal for all connections.

Network architecture. We simulated a microscope focal plane area of L=500 μm^2 as a 2D-layer-like network with periodic boundary conditions and an excitatory/inhibitory neuron ratio of 4:1. Neurons were arranged on a grid of 5 μm steps. Every neuron was sparsely connected with the rest of the network with a connection probability drawn from distance-dependent probability distribution. Each neuron was connected with other network neurons depending on their distance r_{ij}, according to an exponential profile:

$$p_{ij} = A_{syn} * exp^{-(\lambda_{syn} r_{ij})} \quad (4)$$

where A_{syn} is the amplitude of the exponential, varying the number of neurons sampled in the chosen radius r_{ij}. The amplitudes A_{syn} were set to have the number of connected neurons roughly match the distributions described by Katzel et al. 2011, Markram et al. 2015, and Seeman et al. 2018, considering the topological distribution and the excitatory/inhibitory neuron ratio ($A_{exc\text{-}exc}$=14, $A_{exc\text{-}inh}$=24, $A_{inh\text{-}exc}$=14, $A_{inh\text{-}inh}$=24), resulting in a degree of ~75 connections per neuron. The decay controller $\boldsymbol{\lambda}_{syn}$ was set to $\boldsymbol{\lambda}_{exc}$=1.2 and $\boldsymbol{\lambda}_{inh}$=1.5. The seed to the random number generator responsible for the sampling of the connections was systematically varied, producing N=20 different networks, but following the same type of exponentially decaying distributions.

Drive. An external population of Poisson drivers was connected to both the RS and FS populations, with uniform random probability drawn from the half-open interval [0,0.01), with a fixed synaptic weight of 0.5 nS.

Simulator. All simulations were performed using the NEST simulator (Diesmann and Gewaltig 2001) controlled through the PyNN interface (Davison et al. 2009).

Supplementary Figures

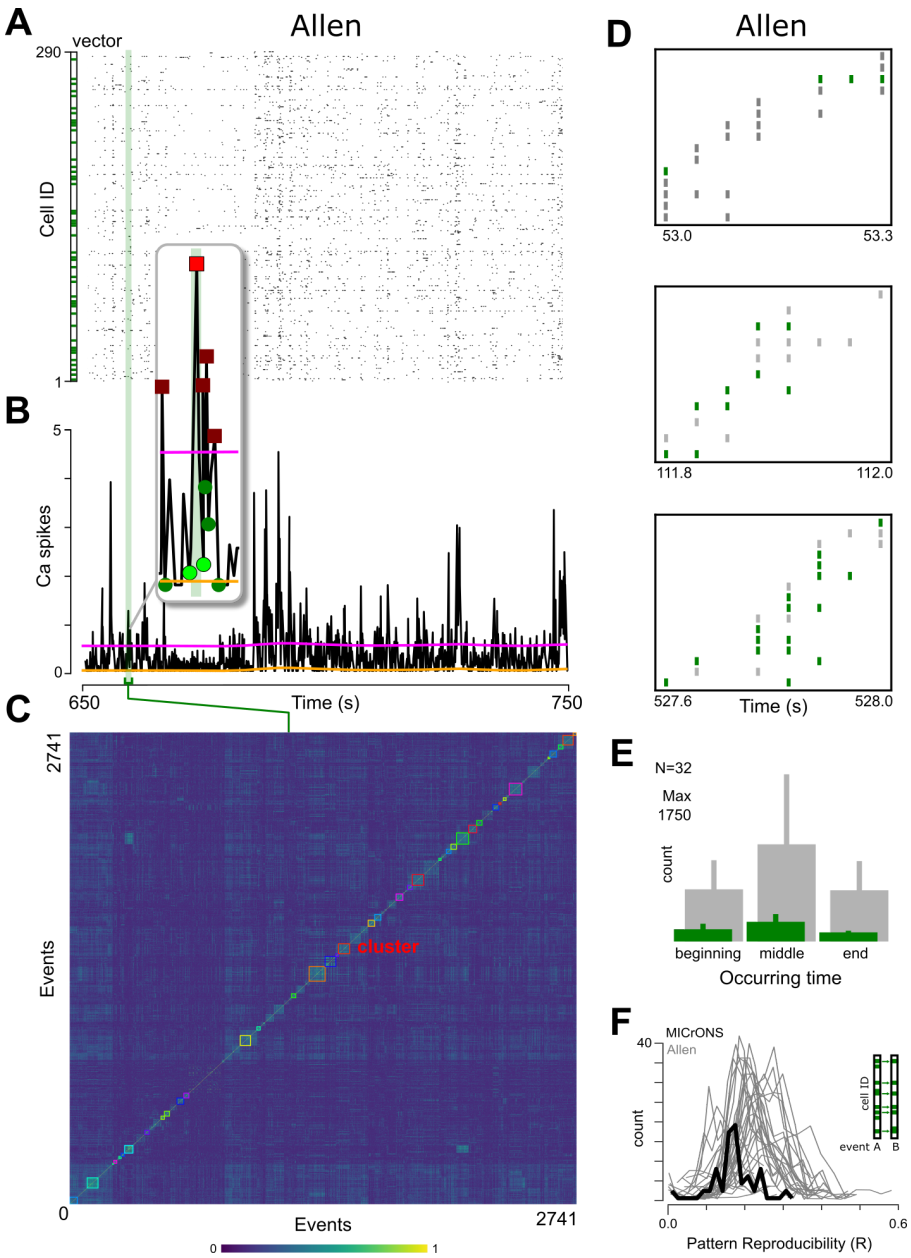

Fig. S1. Comparison of MICrONS and Allen Brain Observatory events. (**A-E**) As for Fig. 1, Fig. 1. (**A**) Calcium spikes raster plot from an example experiment (ID: 540684467). (**B**) Population events. (**C**) Clustered matrix of correlations between population vectors. More events and structure is present in the Allen dataset. However, the stimuli were different (see Allen whitepapers). (**D**) Example events from Allen dataset. (**E**) Core neurons (green) were evenly distributed along events. (**F**) Distributions of pattern reproducibility for the MICrONS phase 1 dataset (black) and each (*N*=32) Allen Brain Observatory (gray). The MICrONS distribution does not qualitatively differ from those from the Allen Brain Observatory. Note that Allen recordings are on average three times longer than the MICrONS, resulting in a reduced count for the latter.

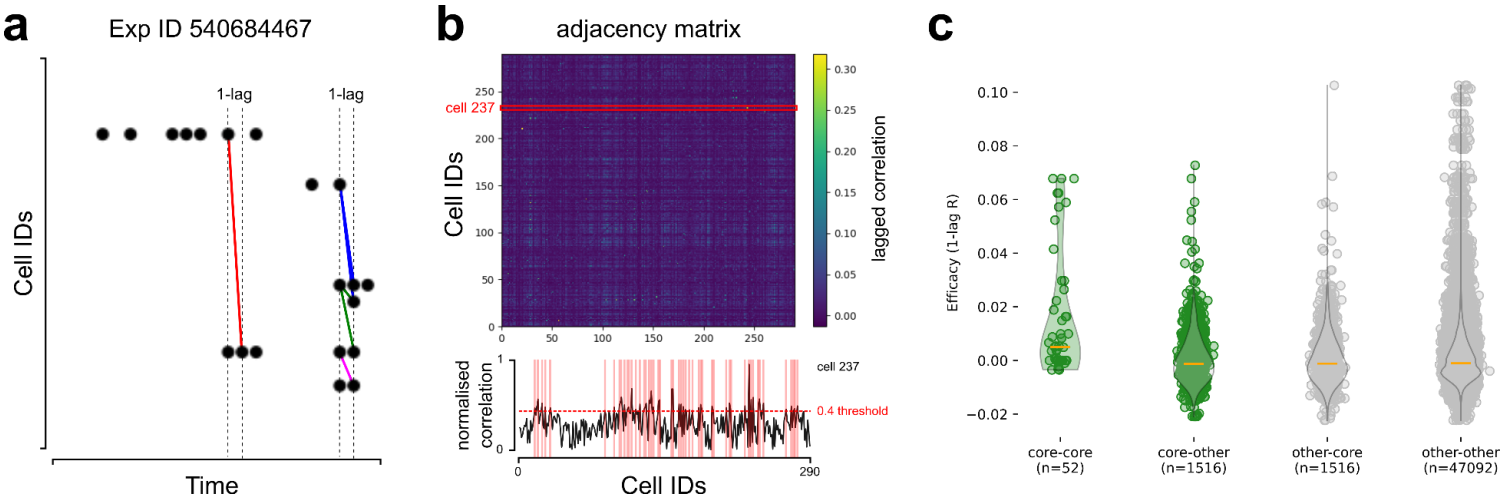

Fig. S2. Core functional correlation. (**A**) Detail from experiment id 540684467 raster plot showing calcium spikes (black circles). Each row contains the calcium spikes for a cell. The minimal interval between two spikes is the two-photon imaging time resolution (30 Hz for the Allen Brain dataset). This interval represents the 1-lag correlation interval to identify cells potentially connected. In the figure, only some cells with 1-lag correlated firing were joined by colored segments to improve legibility. (**B**) The vector of 1-lag correlations for each cell is collected into a matrix (top). Weak correlations (threshold=0.4) were not considered for further analysis (bottom, correlations for cell id 237). (**C**) Core neurons are those sharing (non significantly) more high functional correlation across events compared to others.

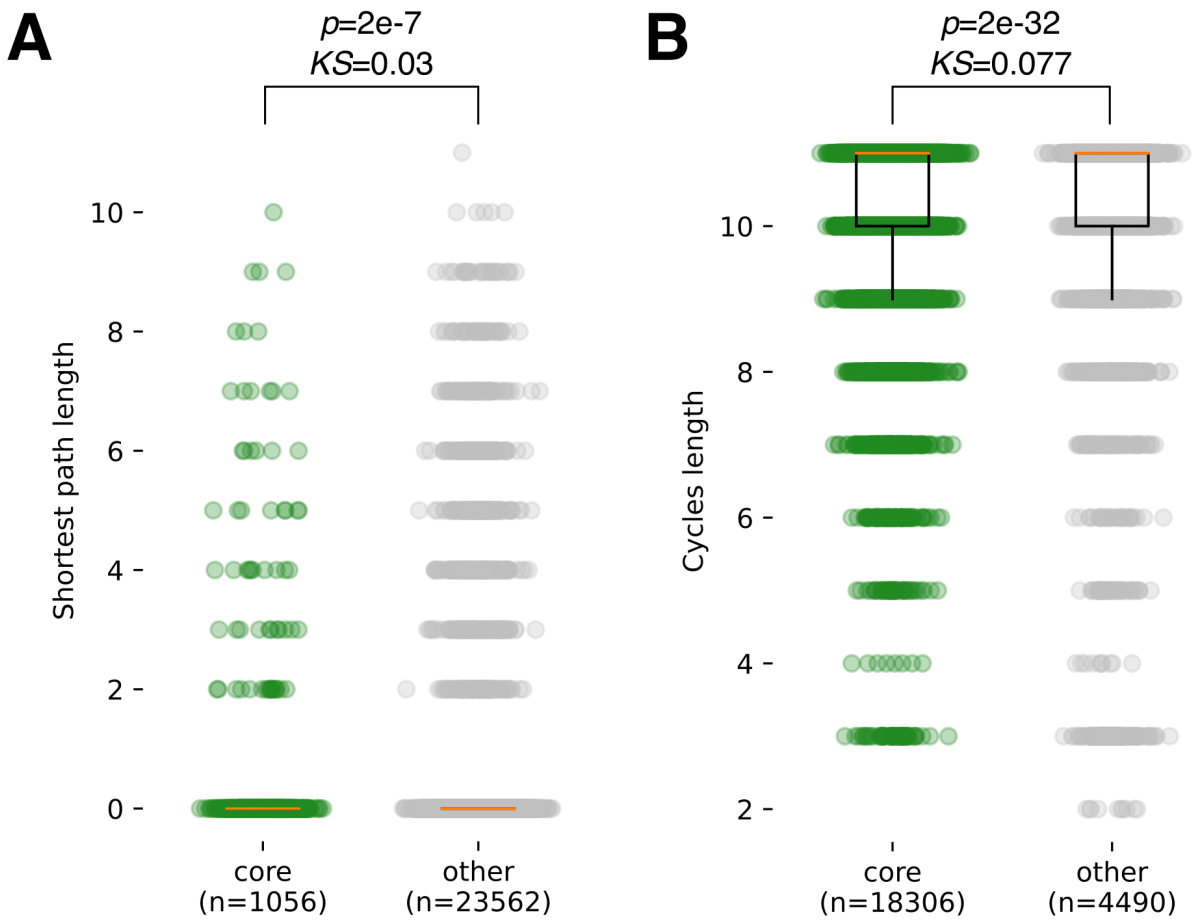

Fig. S3. Cores are not recursively connected via multi-synaptic feedback. Even if cores were not directly connected, they could be highly connected via secondary paths, abundant enough to ensure that they are pattern completion units in traditional attractor network terms. (**A**) Core-to-core shortest paths were neither more nor shorter than other-to-other paths (0.27±1.19 vs 0.16±0.92 length, Kruskal-Wallis t=26.4 p=2.7e-07, *KS*=0.030.) (**B**) Core neurons could be part of looped paths, circling back to them, and providing a weak form of recursion. Core-based cycles were more (green, 18306) compared to other-based cycles (gray, 4490). But core and non-core cycles had on average the same length (10.34±1.07 and 9.93±1.76 connections, with minimal effect size KS=0.077), same as the network diameter (*d*=10).

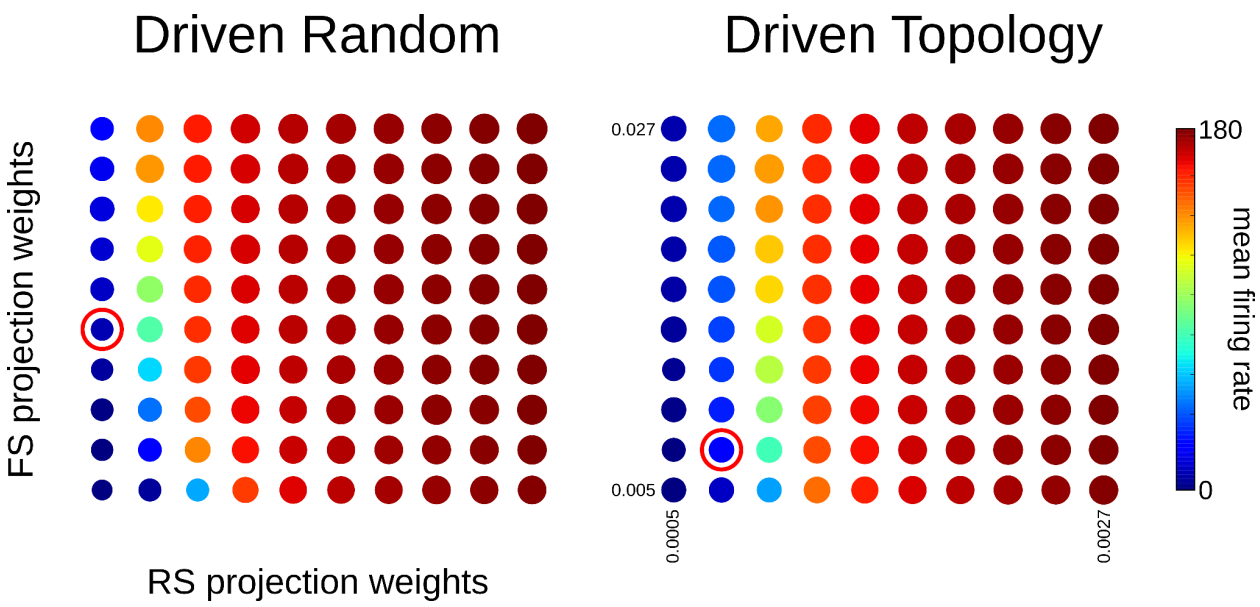

Fig. S4. The model is stable with respect to changes in synaptic parameter space. We explored the parameter space made by the relationship between regular spiking (RS) and fast spiking (FS) synaptic conductance change triggered by incoming spikes, for both the randomly connected (*left*) and topologically connected (*right*), for what concerns the stability of firing regime. Many couples of synaptic parameters produced stable firing rates too high to be comaprable with the 2-photon activities (from 80 spikes/s on, from cyan to red, see colorbar on the right). We therefore further tested several couples below 80 spikes/s and settled on 5-10 spikes/s (marked by the red circle).

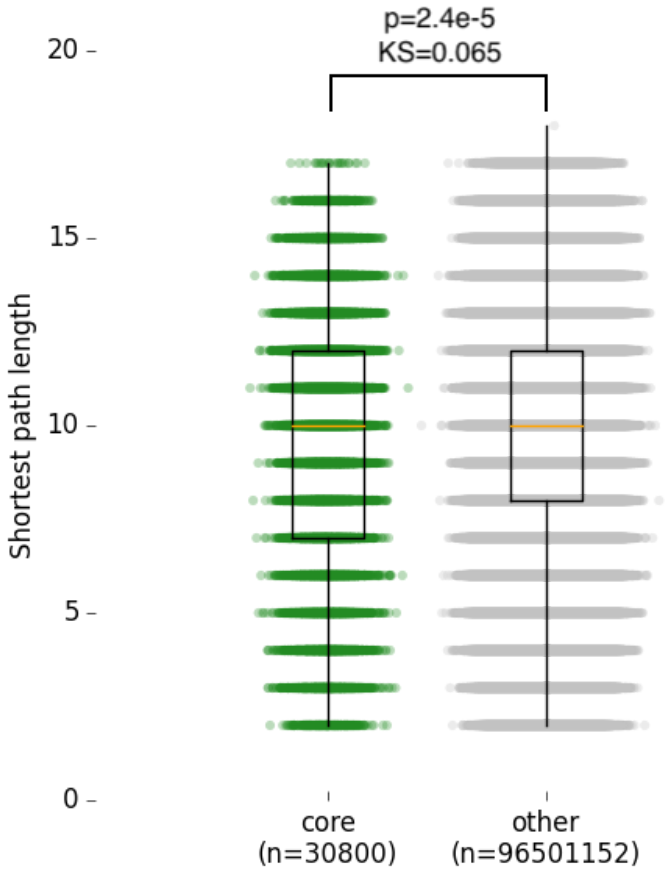

Fig. S5. The shortest paths between model cores and others were not different. Even if cores were not directly connected, they could be highly connected via secondary paths, abundant enough to ensure that they are pattern completion units in traditional attractor network terms (see the works of Amit). To evaluate this aspect also in the model as in the data (see Fig. S2), we measured the number and length of shortest paths between cores and compared them with those made by others. We found that the shortest paths between cores were significantly shorter compared to those between others (cores: 9.75±3.36 vs others: 10.04±2.98 length), overall with negligible effect size (*KS*=0.065). A direct comparison of cycles in the model is computationally unfeasible.

.

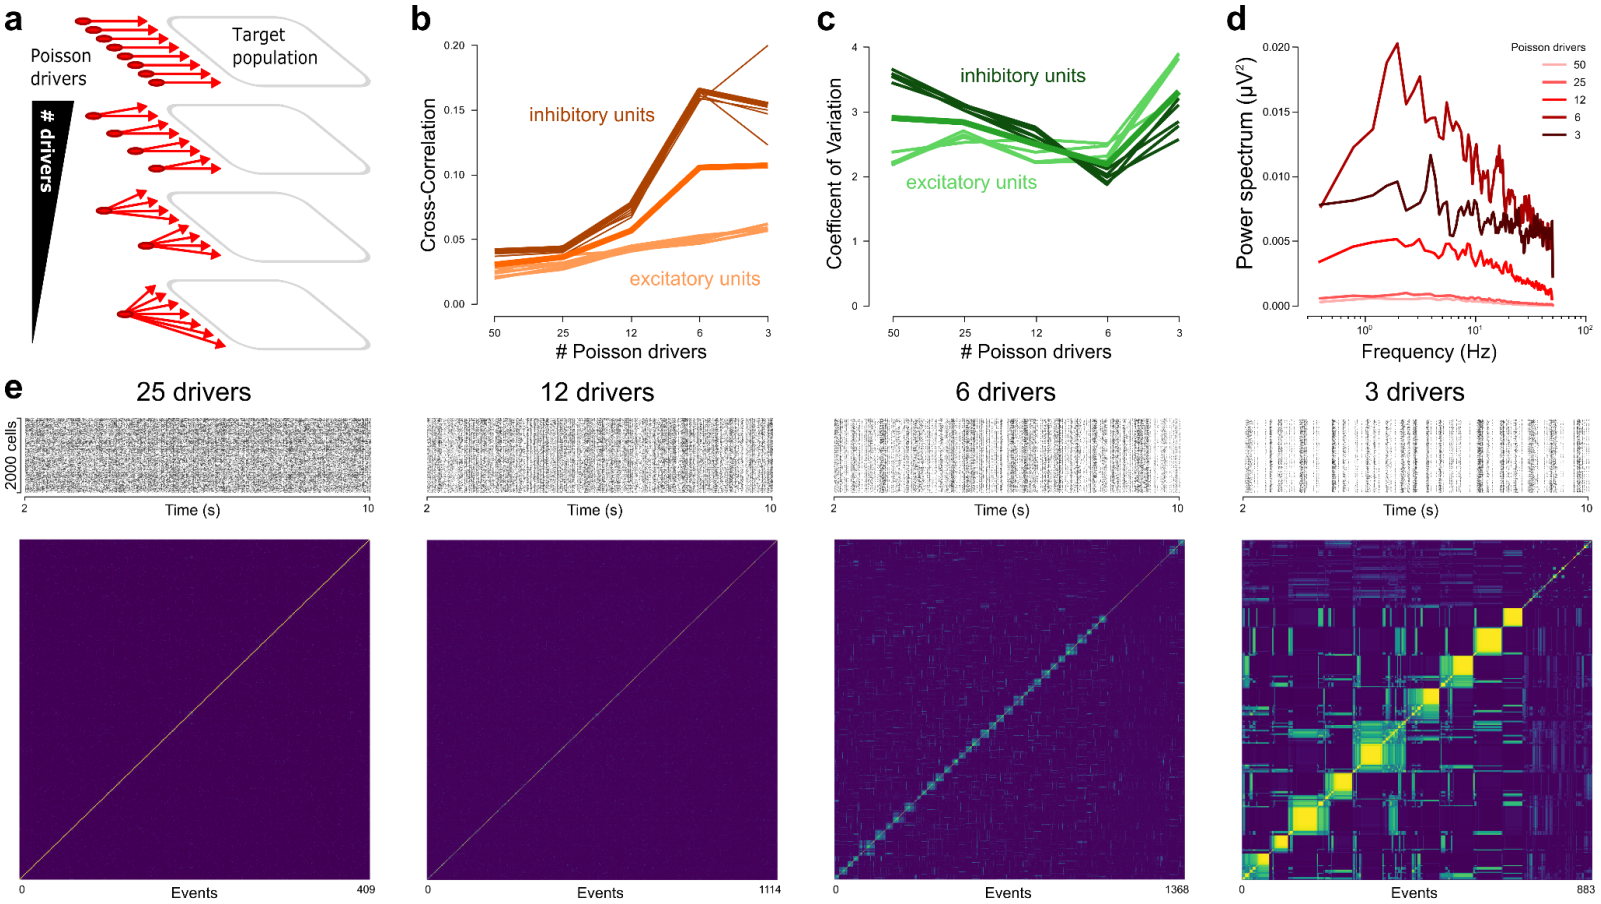

Fig. S6. Increasing input correlations brings the model towards oscillatory regimes. (**a**) For the same target population (right), we systematically reduced the number of Poisson input drivers (red disks and arrows), to explore how the increase in input correlations affected the correlations between population events. (**b**) Reducing the number of Poisson input drivers (n=[50, 25, 12, 6, 3]) increased the population cross-correlation of spikes toward values characteristic of oscillating regimes (from CC=0.035 at n=50, to CC=0.11 at n=3). (**c**) Reducing the number of Poisson input drivers did not significantly affect the population coefficient of variation (Kruskal-Wallis test, p=0.21). (**d**) Reducing the number of Poisson input drivers increased the power for slow oscillations in the Fourier spectrums of the firing rates. (**e**) Four 10 sec spike rasters for 2000 example cells, with their correlation matrices, show the changes in firing regimes responsible for the increase in population events correlation.