



HAL
open science

Perturbed Information Processing Complexity in Experimental Epilepsy

Wesley Clawson, Benjamin Waked, Tanguy Madec, Antoine Ghestem, Pascale
P Quilichini, Demian Battaglia, Christophe Bernard

► **To cite this version:**

Wesley Clawson, Benjamin Waked, Tanguy Madec, Antoine Ghestem, Pascale P Quilichini, et al..
Perturbed Information Processing Complexity in Experimental Epilepsy. *Journal of Neuroscience*,
2023, 43 (38), pp.6573-6587. 10.1523/JNEUROSCI.0383-23.2023 . hal-04295741

HAL Id: hal-04295741

<https://hal.science/hal-04295741>

Submitted on 20 Nov 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Copyright

Perturbed Information Processing Complexity in Experimental Epilepsy

Wesley Clawson,^{1,2} Benjamin Waked,¹ Tanguy Madec,¹ Antoine Ghestem,¹ Pascale P. Quilichini,^{1,*} Demian Battaglia,^{1,3,*} and Christophe Bernard^{1,*}

¹Aix Marseille Université, INSERM, INS, Institut de Neurosciences des Systèmes, Marseille, France, ²Allen Discovery Center, Tufts University, Medford, Massachusetts, and ³University of Strasbourg Institute for Advanced Studies, Strasbourg, France

Comorbidities, such as cognitive deficits, which often accompany epilepsies, constitute a basal state, while seizures are rare and transient events. This suggests that neural dynamics, in particular those supporting cognitive function, are altered in a permanent manner in epilepsy. Here, we test the hypothesis that primitive processes of information processing at the core of cognitive function (i.e., storage and sharing of information) are altered in the hippocampus and the entorhinal cortex in experimental epilepsy in adult, male Wistar rats. We find that information storage and sharing are organized into substates across the stereotypic states of slow and theta oscillations in both epilepsy and control conditions. However, their internal composition and organization through time are disrupted in epilepsy, partially losing brain state selectivity compared with controls, and shifting toward a regimen of disorder. We propose that the alteration of information processing at this algorithmic level of computation, the theoretical intermediate level between structure and function, may be a mechanism behind the emergent and widespread comorbidities associated with epilepsy, and perhaps other disorders.

Key words: algorithm; complexity; entorhinal cortex; epilepsy; hippocampus; information processing

Significance Statement

Comorbidities, such as cognitive deficits, which often accompany epilepsies, constitute a basal state, while seizures are rare and transient events. This suggests that neural dynamics, in particular those supporting cognitive function, are altered in a permanent manner in epilepsy. Here, we show that basic processes of information processing at the core of cognitive function (i.e., storage and sharing of information) are altered in the hippocampus and the entorhinal cortex (two regions involved in memory processes) in experimental epilepsy. Such disruption of information processing at the algorithmic level itself could underlie the general performance impairments in epilepsy.

Received Mar. 3, 2023; revised July 7, 2023; accepted July 20, 2023.

Author contributions: W.C., A.G., P.P.Q., D.B., and C.B. designed research; W.C., A.G., and P.P.Q. performed research; W.C., B.W., T.M., P.P.Q., and D.B. analyzed data; W.C., B.W., T.M., P.P.Q., D.B., and C.B. wrote the first draft of the paper; W.C., A.G., P.P.Q., D.B., and C.B. edited the paper; W.C., P.P.Q., D.B., and C.B. wrote the paper.

W.C. was supported by the M-GATE program, which has received funding from the European Union's Horizon 2020 research and innovation program under Marie Skłodowska-Curie Grant Agreement 765549. T.M. was supported by Aix-Marseille Université. P.P.Q. was supported by FRM, FFRE and CURE Epilepsy Taking Flight Award. D.B. was supported by French Agence Nationale pour la Recherche ERMUNDY ANR-18-CE37-0014-02 and University of Strasbourg Institute for Advanced Study for a Fellowship, within the French national program "Investment for the future." C.B. was supported by ANR 19-CE14-0036-01. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript. We thank Romain Goutagny and Anna Levina for scientific discussions and comments regarding this work.

*P.P.Q., D.B., and C.B. contributed equally to this work.

The authors declare no competing financial interests.

Correspondence should be addressed to Wesley Clawson at wesley.clawson@tufts.edu.

<https://doi.org/10.1523/JNEUROSCI.0383-23.2023>

Copyright © 2023 the authors

Introduction

Most, if not all, neurologic pathologies, including epilepsies, Alzheimer's disease, and Parkinson's disease, display commonalities in terms of cognitive (e.g., memory) and psychiatric (e.g., anxiety and depression) comorbidities (Hesdorffer, 2016). Most mechanistic based studies correlate high-level to "structure level" alterations, that is, gene/protein expression, anatomy (cell death, sprouting), or neuronal properties (Berchtold et al., 2014; Cho et al., 2015; Holmes, 2015). However, changes at the structural level are often very different from one pathology to the next and may vary even within a specific pathology (Gorur-Shandilya et al., 2022), in keeping with the principle of degeneracy (Edelman and Gally, 2001; Mason, 2015). If different network architectures can produce the same physiological activity (Marder and Goaillard, 2006; Goaillard and Marder, 2021), different network alterations may produce the same deficits. However, degeneracy may not be present at a level higher than the structural one. Distinct structural alterations could produce similar changes in primitive information processing operations (the algorithmic level), providing a

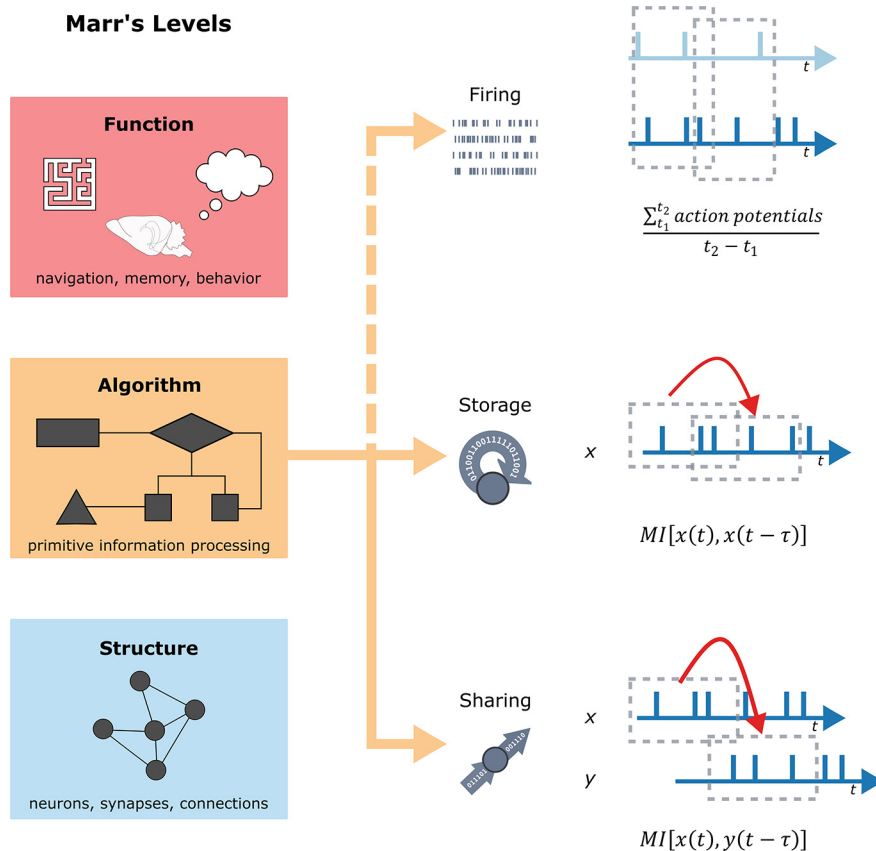


Figure 1. Marr's levels and analytical approach. Diagram example of Marr's levels. Each level is important in and of itself; however, the algorithmic level is often not considered. We interpret the algorithmic level as encompassing very simple operations on the information conveyed by spikes which directly stem from patterns of coordinated firing. Neuronal firing is not arbitrary but constrained by "rules of engagement" (shaped on their turn by the underlying structural level), that is, rules that govern how neurons fire individually and as part of a group. Because of these rules, the information carried by firing (top diagram) undergoes a basic processing; for example, it is actively maintained through time by some neurons (AIS, middle diagram) or it is copied to the spiking of other neurons (active information sharing, bottom diagram). Both firing and storage are measures of single-neuron spiking activity (x), while sharing is a measurement taken across the entire network (x and y).

framework to explain why cognitive comorbidities are shared within and across neurologic disorders. The primitive information processing operations are basic and "prefunctional" but necessary to any cognition-relevant neural computation (Lizier, 2013; Taatgen, 2013; Voges et al., 2022). Whether primitive information processing operations are altered in neurologic disorders is not known. Here we study information storage and sharing, as basic information processing operations, using temporal lobe epilepsy (TLE) as a model of neurologic disorder.

TLE is the most common form of epilepsy in adults (Tatum, 2012); and it is highly heterogeneous in terms of histopathology (Blümcke et al., 2013), semiology (Barba et al., 2007; Bartolomei et al., 2008), and cognitive and psychiatric comorbidities (Holmes, 2015; de Barros Lourenco et al., 2020; Krishnan, 2020). Such heterogeneity is also found in experimental models of TLE (Rusina et al., 2021). Structural alterations may change several features that are relevant for any kind of information processing, such as rate coding, temporal coding, synaptic plasticity, and network oscillations (Lenck-Santini and Scott, 2015). Many phenomena considered fundamental for functional computations are altered in experimental TLE: hippocampal place cells are unstable, firing becomes randomized during ripples, and oscillatory activities are modified, in correlation with cognitive deficits (Lenck-Santini and Holmes, 2008; Chauvière et al., 2009; Suarez et al., 2012; Inostroza et al., 2013; Lopez-Pigozzi et al., 2016; Valero et al., 2017). However, the mechanistic link between structural and

functional alterations in TLE is missing. Disruptions of primitive information processing may provide such a link.

To gain access to the level intermediate between structure and function (Fig. 1, left), denoted as algorithmic by Marr and Poggio (1977), we begin from the tenet that neural information is conveyed by action potentials and that one form of processing stems from spatiotemporally organized neuronal firing. Neuronal firing carries information because of its statistical properties, and information theory (Shannon, 1948) established ways to measure information in a data-driven manner agnostic to the content itself. Using these tools, we study two primitive information processes: information storage and information sharing (Fig. 1, right). For a single neuron, auto-correlation in firing maintains and buffers information through time, an operation named active information storage (AIS) (Lizier et al., 2012; Wibral et al., 2014). Cross-correlated firing between neurons allows for information sharing between them (Kirst et al., 2016). Focusing on such basic operations gives access to the algorithmic level (Clawson et al., 2019).

As neural computation is brain state-dependent (Quilichini and Bernard, 2012), we consider the global brain states of theta (THE) and slow oscillations (SOs), which can be easily recorded during anesthesia. Previous work in control animals demonstrated that neuronal activity patterns in the hippocampus (HPC) and entorhinal cortex switch between different information processing substates (IPSS) (Clawson et al., 2019). An IPS corresponds to an epoch in which primitive operations of information storage and sharing in a local microcircuit remain

temporally consistent. IPSs continuously switch from one IPS to another, similarly to what has been described at higher level of organization, such as the dynamics of resting state networks and EEG microstates (Van de Ville et al., 2010; Calhoun et al., 2014). In the control HPC and entorhinal cortex, the sequences of IPSs are complex (i.e., standing between order and disorder) (Clawson et al., 2019).

Here, we reintroduce the notion of IPS; and using an unbiased quantification we compare their properties and organization between control and experimental TLE conditions. We focus on the HPC and the entorhinal cortex, two major structures commonly affected in TLE (Curia et al., 2008). We find that IPS's internal organization and switching dynamics are preserved in epilepsy but are shifted toward a less structured and more random spatiotemporal organization in experimental epilepsy than in control. Such disruption of information processing at the algorithmic level itself could underlie the general performance impairments in TLE.

Materials and Methods

Ethics. All experiments were conducted in accordance with Aix-Marseille Université and Institut National de la Santé et de la Recherche Médicale Institutional Animal Care and Use Committee guidelines. The protocol was approved by the French Ministry of National Education, Superior Teaching, and Research under authorization 01451-02. All surgical procedures were performed under anesthesia, and every effort was made to minimize suffering and maximize the animals' well-being from their arrival to their death. All the animals were housed in pairs in large cages with minimal enrichment, food and water at libitum, in a room with controlled environment (temperature: $22 \pm 1^\circ\text{C}$; 12 h light/dark schedule with lights off at 8:00 P.M.; hygrometry: 55%; ventilation: 15–20 vol/h).

Experimental design and statistical analysis. We use in this work a portion of the data (5 of 7 original experiments) initially published by Clawson et al. (2019) as control data, which includes local field potentials (LFPs) and single-unit recordings obtained from the dorsomedial entorhinal cortex (mEC) and the dorsal HPC of anesthetized rats. The remaining 2 were not included as controls in this study as they were recordings in the mPFC, a region not explored directly in the new recordings in epileptic rats. Six recordings are original data, which includes LFPs and single units recorded in the mEC and HPC recorded simultaneously under anesthesia in epileptic condition. These sample sizes were not predetermined, and all available data that fit the region criteria and number of recorded neurons >20 were used. For details on individual recordings and number of cells, see Extended Data Figure 2-1.

Statistical tests, where performed, were done as two-sided permutation t tests. Their values are reported in Extended Data Figures 2-3 and 2-6. In all cases, we follow the statistical analysis detailed in Ho et al. (2019). Captions were written with the method used, but often distributions were compared using distributions of mean difference with the following text taken from Ho et al. (2019):

“The P value(s) reported are the likelihood(s) of observing the effect size(s) if the null hypothesis of zero difference is true. For each permutation P value, 5000 reshuffles of the control and test labels were performed. They are included here to satisfy a common requirement of scientific journals.”

Additionally, throughout the text lines as a function of k_{tot} are shown with the bold line representing the mean and the shaded, colored area around representing a 99% CI.

Epilepsy model and surgery. We induced status epilepticus (SE) on 6 male Wistar (250–400 g; Charles Rivers) by a single intraperitoneal injection of pilocarpine (320 mg/kg; Sigma-Aldrich), 1 week after receiving the animals from the vendor. To reduce peripheral effects, rats were pretreated with methyl-scopolamine (1 mg/kg, i.p.; Sigma-Aldrich) 30 min before the pilocarpine injection. SE was stopped by diazepam (10 mg/kg, i.p., two doses within a 15 min interval) after 60 min. Then the animals were

hydrated with saline (2 ml, i.p., twice within 2 h) and fed with a porridge made of soaked pellets, until they resumed normal feeding behavior.

At least 8 weeks after SE induction, we performed acute recordings. Only rats that had undergone SE and in which spontaneous seizures were subsequently observed were used for the recordings. Rats were first quickly placed in isoflurane (4% in 2 L/min O_2) and injected intraperitoneally with urethane (1.5 g/kg) and ketamine/xylazine (20 and 2 mg/kg, i.m.), additional doses of ketamine/xylazine (2 and 0.2 mg/kg) being supplemented during the electrophysiological recordings. At all times, the body temperature was monitored and kept constant with a heating pad. Heart rate, breathing rate, pulse distension, and arterial oxygen saturation were also monitored with an oximeter (MouseOX; StarrLife Sciences) during the duration of the experiment to ensure the stability of the anesthesia and monitor the vital constants. The head was fixed in a stereotaxic frame (Kopf) and the skull was exposed and cleaned. Two miniature stainless-steel screws driven into the skull above the cerebellum served as ground and reference electrodes. Two craniotomies were performed to reach the mEC and the CA1 field of the HPC, respectively: from bregma: -7.0 mm AP and 4.0 mm ML; and from bregma: -3.0 mm AP and 2.5 mm ML. We chose these coordinates to respect known anatomic and functional connectivity in the cortico-hippocampal circuitry (Witter et al., 1988, 1989). Two 32-site silicon probes (NeuroNexus) were mounted on a stereotaxic arm each. A H1x32-10 mm-50-177 was lowered at 5.0–5.2 mm from the brain surface with a 20° angle to reach the dorsomedial portion of the mEC, and a H4x8-5 mm-50-200-177 probe was lowered at 2.5 mm from the brain surface with a 20° angle to reach dorsal CA1. The online positioning of the probes was assisted by the following: the presence of unit activity in cell body layers and the reversal of THE ($[3, 6]$ Hz in anesthesia) oscillations when passing from layer 2-1 for the mEC probe, and the presence in stratum pyramidale either of unit activity and ripples (80–150 Hz) for the HPC probe. At the end of the recording, the animals were injected with a lethal dose of pentobarbital Na (150 mg/kg, i.p.) and perfused intracardially with 4% PFA solution. We confirmed that the position of the electrodes (DiIC18(3), catalog #46804A, InterChim) was applied on the back of the probe before insertion) histologically on $40 \mu\text{m}$ Nissl-stained section as reported previously in detail (Quilichini et al., 2010; Ferraris et al., 2018). We used only experiments with appropriate position of the probe for analysis.

Data collection and spike sorting. Extracellular signal recorded from the silicon probes was amplified ($1000\times$), bandpass filtered (1 Hz to 5 kHz), and acquired continuously at 32 kHz with a 64-channel DigitalLynx (NeuraLynx) at 16-bit resolution. We preprocessed the raw data using a custom-developed suite of programs (Csicsvari et al., 1999). The signals were downsampled to 1250 Hz for the LFP analysis. Spike sorting was performed automatically, using KLUSTAKWIK (<http://klustakwik.sourceforge.net>) (Harris et al., 2000), followed by manual adjustment of the clusters, with the help of auto-correlogram, cross-correlogram, and spike waveform similarity matrix (KLUSTERS software package, <http://klusters.sourceforge.net>) (Hazan et al., 2006). We did not filter on neuronal type, such as excitatory or inhibitory, and all neurons were included in analysis with an average of 53 neurons per control animal and 60 per epileptic animal (Extended Data Fig. 2-1). After spike sorting, we plotted the spike features of units as a function of time, and we discarded the units with signs of significant drift over the period of recording. Moreover, only units with clear refractory periods and well-defined cluster were included in the analyses (Harris et al., 2000). Recording sessions were divided into brain states of THE and SO periods using a visual selection from the ratios of the whitened power in the HPC LFP $[3, 6]$ Hz THE band and the power of the mEC LFP neighboring bands ($[1, 3]$ Hz and $[7, 14]$ Hz), and assisted by visual inspection of the raw traces (Quilichini et al., 2010; Ferraris et al., 2018). We then used band-averaged powers over the same frequency ranges of interest as features for the automated extraction of spectral states via unsupervised clustering, which confirmed our manual classification. We determined the layer assignment of the neurons from the approximate location of their soma relative to the recording sites (with the largest-amplitude unit corresponding to the putative location of the soma), the known distances between the recording sites, and the histologic reconstruction of

the recording electrode tracks. Animals were recorded for at least 2 h to get few alternations of THE and SO episodes.

Feature computation. As in our previous work, for each region recorded, we computed four main features from the electrophysiological data: global oscillatory band, neuronal firing sets, AIS, and the information sharing. We also keep the same sliding window paradigm where each feature is computed within a 10 s window, and then the window is moved forward in time 1 s, which gives a 9 s overlap. These values for windows and overlap were chosen for two reasons. First, we chose a long window as measures of mutual information (MI) measures require many samples to properly estimate probabilities involved. In the slowest global state (SO), a 10 s window would capture ~10 cycles, which seemed reasonable for estimation. Second, we chose an overlap of 9 s, meaning that there is a 1 s difference between two temporally adjacent windows, or one SO cycle (more for THE). Given that replay events, relatively fast known phenomenon, occur within 500 ms (Lee and Wilson, 2002), we posit that double this would allow for both the capturing of fast events while allowing for conceptually clear measures of correlation in the population. With a 1 s shift, the correlation of a given feature between two temporally adjacent windows highlights how the network can shift on a second timescale. Overall, when features are computed as described below, they are computed in this windowed fashion. The global oscillatory band features were computed by examining the LFP from both EC and CA1 and computing spectral power within 8 unequally sized frequency ranges (0–1.5, 1.5–2, 2–3, 3–5, 5–7, 7–10, 10–23, and 23–50 Hz), averaged over all channels within each of the recorded layers.

Firing sets, AIS, and the information sharing networks were all computed using a binarized raster built from the temporal labeling of spike firing (see Data collection and spike sorting). Spiking data were binned using a 50 ms bin; if a neuron fired within a given bin, the output is a 1, and if not, a 0. This, for example, would mean that a 2 h recording would be transformed from a $7200 \text{ s} \times N$ neuron matrix to a $7,200,000 \times N$ neuron matrix that is composed solely of 0's and 1's. Firing sets were computed by computing the average firing density for each neuron within a window, and after these averages were compiled into time-dependent vectors. This resulting matrix is the Firing Features. AIS was computed by measuring the MI of a neuron's binarized spike train between a given window and the window previous. What AIS seeks to capture is the temporal ordering of individual spiking neurons, rather than capturing neurons that fire temporally close to one another (e.g., in the firing features). The resulting matrix is the Storage Features. AIS is meant to quantify how much the activity of a unit is maintaining over time information that it was conveying already in the past (Lizier, 2013). It is an activity-based metric (hence the adjective "active"), able to detect when temporal patterns in the activity of a single unit can serve the functional role of "memory buffer." AIS is strictly defined as follows:

$$\text{AIS}_i = \text{MI}[i(t), i(_t)]$$

i.e., as shared information between the present activity $i(t)$ of a single unit i and its past history of activity $i(_t)$. Before computing MI, we binned all spike trains with method as for determining the Firing(t) descriptive feature vector. The limited amount of available data within each temporal window makes it necessary to introduce approximations. Therefore, we replaced the full past history of activity $i(_t)$ with activity at a time in the past $i(t - \tau)$ and then summed over all the possible lags as follows:

$$\hat{\text{AIS}}_i = \sum_{\tau} \text{MI}[i(t), i(t - \tau)]$$

where the lag τ was varying in the range $0 \leq \tau \leq 0.5 T_{\theta}$, where T_{θ} is the phase of the theta cycle. MI values were generally vanishing for longer latencies. We evaluated MI terms using a "plug-in" function estimator on binarized spike trains, which takes the binned spike trains of two neurons for a defined time window and computes the MI and entropy values of the two variables (Lizier, 2013). We estimated the probability p that a bin includes a spike and the complementary probability $1 - p$ that a bin

is silent for each unit, by direct counting of the frequency of occurrence of 1's and 0's in the binned spike trains of each unit. These counts yielded the probability distributions $P(I)$ and $P(j)$ that two neurons i and j fire or not. Analogously, we sampled directly from data the histogram $P(i, j)$ of joint spike counts for any pair of two units i and j . These histograms were then directly "plugged in" (hence the name of the used estimator) into the definition of MI itself as follows:

$$\text{MI}(i, j) = \sum_i \sum_j P(i, j) \log_2 \frac{P(i, j)}{P(i)P(j)}$$

We then subtracted from each MI value a significance threshold (95th percentile of MI estimated on shuffled binarized trains, 1000 replicas), putting to zero nonsignificant terms (and thus negative after bias subtraction).

Information sharing is computed by measuring the MI between a given neuron's binarized spike train within a window and another neuron's binarized spike train in the window previous. This process is iterated over all possible neuron pairs. Information sharing captures a similar metric to that of AIS, although the key difference is that information sharing captures not just the temporal ordering, but the spatiotemporal ordering of spike timing, as it is computed across neuron pairs, rather than individual neurons. The resulting matrix is the Information Sharing. Within each time window, we computed time-lagged $\text{MI}[i(t), j(t - \tau)]$ between all pairs of spike density time series for different single units i and j . Although MI is not a directed measure, a pseudo-direction of sharing is introduced by the positive time lag, supposing that information cannot be causally shared from the future. Thus, for every directed pair of single units i and j (including auto-interactions, with $i = j$), we defined pseudo-directed information sharing as follows:

$$I_{\text{shared}}(j \rightarrow i) = \sum_{\tau} \text{MI}[i(t), j(t - \tau)]$$

where the lag τ was varying in the range $0 \leq \tau \leq 0.5 T_{\theta}$, where T_{θ} is the phase of the theta cycle. Once again, we estimated MI terms via direct plug-in estimators on binarized spike trains, as with storage, subtracting a significance threshold (95th percentile of MI estimated on shuffled binarized trains, 400 replicas) and zeroing out significant terms. All these $I_{\text{shared}}(j \rightarrow i)$ entries were interpreted as weights in the adjacency matrix of an information sharing directed functional network, and we defined as sharing assembly formed by a neuron i the star-subgraph of the information sharing network composed of i and all its immediate neighbors. We compiled all the overall N^2 different values of $I_{\text{shared}}(j \rightarrow i)$ into time-dependent feature vectors $\text{Sharing}_A(t)$, describing thus all the possible sharing assemblies at a given time. Although these measures have only been briefly described here, we suggest to the interested reader to examine the methods presented in our previous work where they have been rigorously defined and discussed (Clawson et al., 2019).

Feature-based substate extraction. State extraction for each recording were also computed using the methods of our previous work, namely, based around k -means clustering of each feature. The exception here is we no longer choose a stable number of K clusters in k -means. Rather, we cluster our three raster-based computed features (firing, storage, sharing) three separate times with K ranging from $K = 3, 4, \dots, 10$. The function "kmeans" was used from the default MATLAB toolbox. More information can be found on The MathWorks website. These K values were chosen as they represented a clustering range of too gross to too fine based on previous findings. $K \leq 2$ would represent the same, or less, number of states as global states, which was previously established to be too small (Clawson et al., 2019). The clustering became too fine when $K \geq 10$, wherein many substates only appeared for brief time periods, and never re-occurred. For each feature, there are 8 different clustering results, done in an unsupervised manner 3 times to ensure that our results do not rely on single instance of clustering. This gave our analysis an opportunity to compute all metrics defined below over a robust range of K , ensuring that we can investigate how our substate stable metrics and results vary with arbitrarily too little or too many substates.

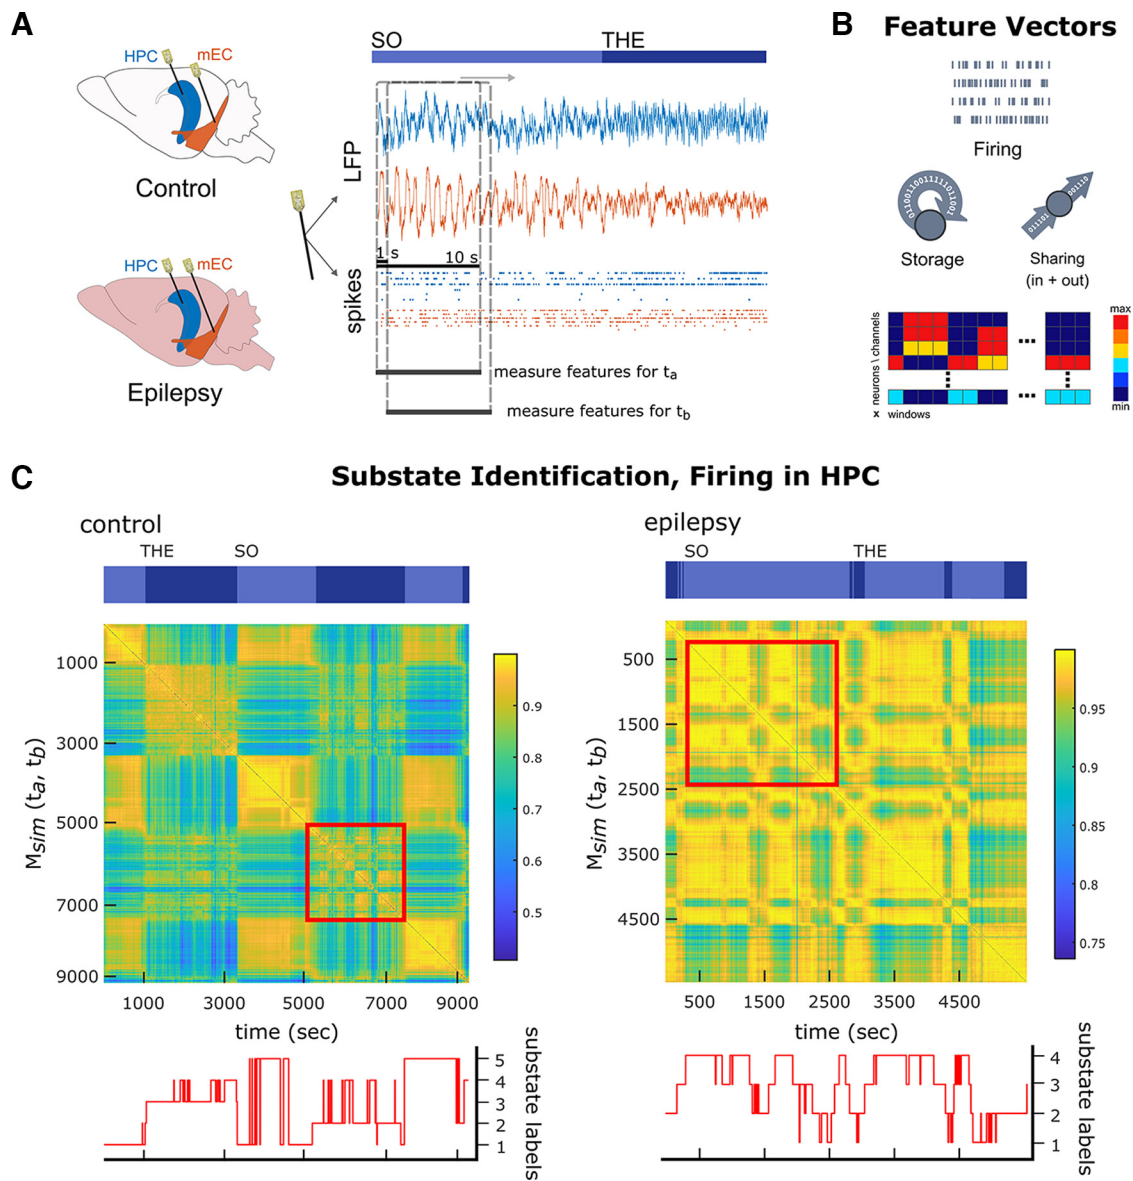


Figure 2. Experimental and analytical design. **A**, Diagram represents the approximate recording locations in mEC (orange) and CA1 (blue) in control and experimental epilepsy. Example of LFP (top) and firing (bottom) [each line indicates one neuron; dot represents an action potential] data recorded in control CA1 and mEC during SO and THE. Overlay represents our analytic method that uses 10-s-long sliding windows shifted by 1 s at each step. **B**, Diagram examples of the three data features shown in Figure 1. **C**, Similarity matrices for CA1 in example control and epileptic animals (see Materials and Methods). Above each matrix is a representation of THE and SO epochs through time. Each point on the similarity matrix represents the correlation of features across all recorded neurons between two time points in the recording. Therefore, large “blocks” of yellow represent an epoch of time where the behavior of the recorded neurons, or network, is very similar and stable (i.e., a substate). Blocks of blue represent when the population activity is very uncorrelated between two time points. In both conditions, substate activity can be seen by eye. An example: in the control condition from time ~5200 to ~7300 s, the red square highlights a period of stable THE oscillations, yet contains many small blocks, or reoccurring substates. Below each matrix is an example of clustering to identify substates through time (5 substates in the control condition and 4 in the epileptic condition). While some substates appear to be “blips,” this is because of the size of the recording — these small substates last 10–20 s (70–140 theta cycles). The same can be seen in the red square in the epileptic condition from time ~500 to ~2500 s. While the figures in **C** are presented next to one another, they are not paired in any way, as they are recordings of 2 different animals. For number of neurons recorded, see Extended Data Figure 2-1. For additional details on spectral analysis between control and epilepsy and the related statistics, see Extended Data Figures 2-2, 2-3, 2-4, 2-5, and 2-6.

To compute the null model for substate extraction, the process detailed above was repeated with the time stamps of all firing, storage, and sharing jittered. This therefore retains the global mean and variance. Then, *k*-means was run on this jittered dataset 3 times, to produce three different clustering of the randomized dataset. These were not modified after this step and were used in any instances where a null model was needed (i.e., for silhouette and contrast).

Substate tables. Our main meta-object of study is a state table, a combination of our four features into a matrix (4 × number of windows). Table generation is an iterative process, as we have 8 possible substate configurations per feature. First, *k* = 3 in cluster attempt 1 for firing (FIRE_{K3C1}), *k* = 3 in cluster attempt 1 for storage (STORE_{K3C1}), and

k = 3 in cluster attempt 1 for sharing (SHARE_{K3C1}), are used in conjunction with the clustered spectral substates to form substate table 1 (see Fig. 2A).

Then, FIRE_{K3C1}, STORE_{K3C1}, and SHARE_{K4C1} are used in conjunction with the clustered spectral substates to form substate table 2. After, FIRE_{K3C1}, STORE_{K3C1}, and SHARE_{K5C1} used in conjunction with the clustered spectral substates to form substate table 3. This process continues such that all combinations of possible *k* values have been saved for a total of 512 different substate tables, with the final table having FIRE_{K10C1}, STORE_{K10C1}, and SHARE_{K10C1}. It is important to note that all tables have the same spectral clustering, as the 2 substates of SO and THE are extremely robust as discussed above. This entire process is

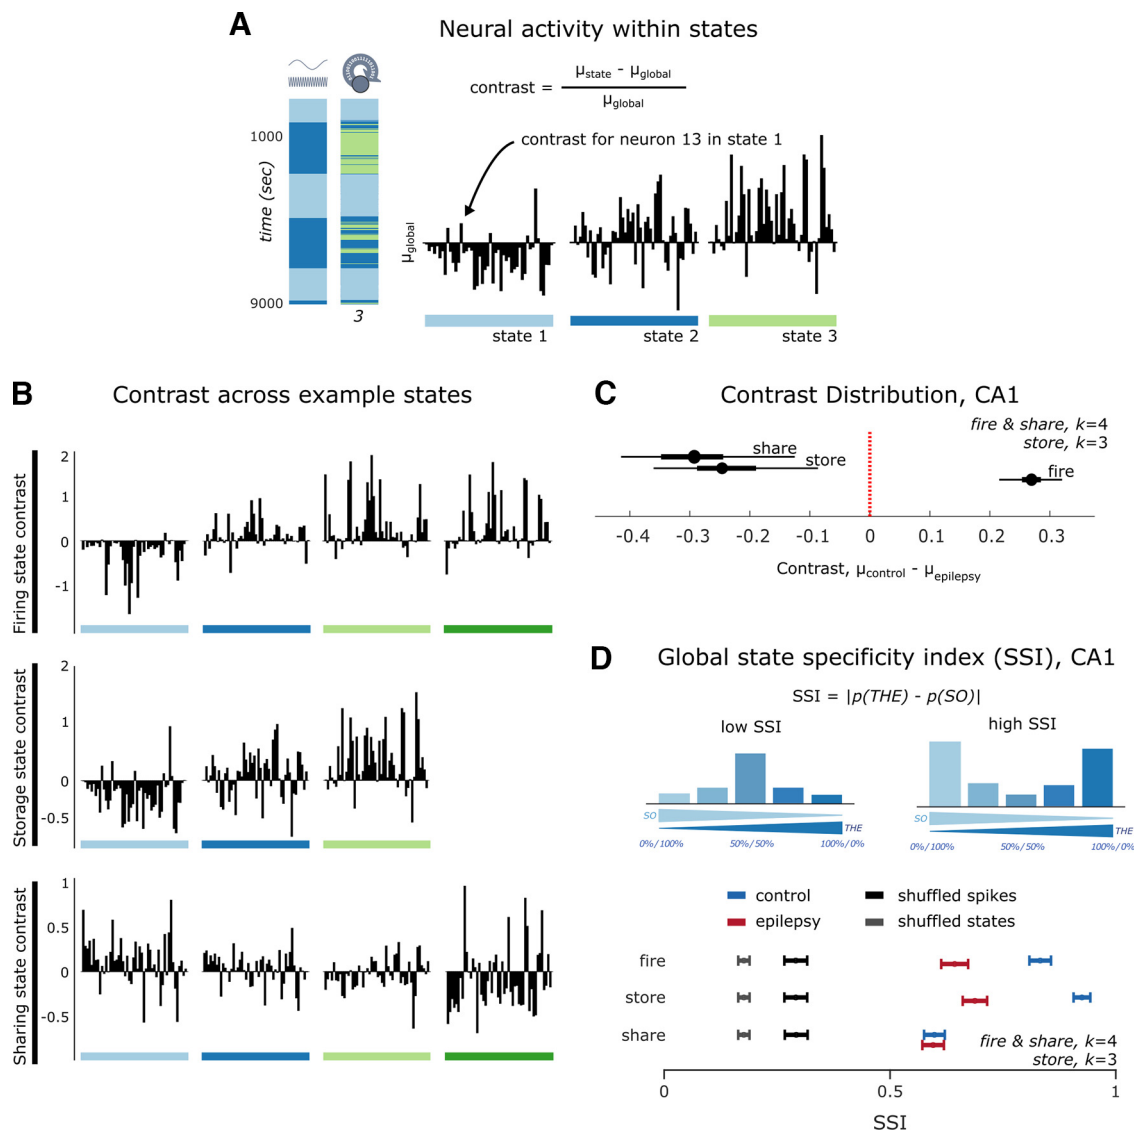


Figure 3. Substate characterization in control and epilepsy. **A**, An enlarged example of contrast in substates. Black bars represent the storage contrast of 55 neurons in the three detected states in an example recording. Left, Spectral states of THE and SO next to the substates of storage. Time goes from top to bottom in this visualization. **B**, A full example of contrast across firing, storage, and sharing features for the same example recording as in **A**. For each, the 0 represents the mean global value for the feature and the height of the bar, the contrast. The “landscape” of these bars highlights the differing activity within each state that the blocks were capturing in Figure 2 via correlations. **C**, The difference of mean distributions between substates from control and epilepsy animals (same states as in **B**; see Materials and Methods). Circles represent the mean difference. Thick black bars represent the 25%–75% quantile. Thin black bars represent the 1%–99% quantile. Red dotted line is to add the null hypothesis line of no significant difference between control and epilepsy. **D**, Visual examples of the two extremes of the measure of SSI. Bottom, The distribution of mean SSI for control (blue), epilepsy (red), and shuffled (black/gray) recordings for firing, storage, and sharing for their respective k . Error bars indicate the 99% bootstrapped CI. For details on shuffling, see Materials and Methods. For more details on the contrast, SSI, and clustering, see Extended Data Figures 3-1, 3-2, 3-3, and 3-4.

then repeated for each clustering attempt, resulting in 3 sets of our 512 substate tables for each region for each recording. Where applicable, all results are given as a function of total k states per table (i.e., for state table 1, there are 2 global states, 3 firing, 3 storage, and 3 sharing, for a total $k_{\text{total}} = 11$).

To produce the ordered tables for the “ordered” null model, each substate table was sorted such that all substates with label 1 appeared first, label 2 was second, and so on. This can easily be achieved with the MATLAB function *sort*. There is only one possible version of this type of ordering; therefore, the sample size for ordered tables is the same as recordings ($n = 5$ for control, $n = 6$ for epilepsy). To produce the randomized tables, substate labels were randomly permuted in time. For this process, we used bootstrapping to produce as 5000 randomizations to ensure the random null model was as strong as possible. To do this, 90% of each table was taken, randomly permuted, and saved. These resulting tables were used as the random null model for relative dictionary and complexity seen in Figures 5 and 6.

Shuffling. To create shuffled datasets of recorded data, time shuffling was used. For spikes, this was done in two ways. The first method for the binarized data was shuffling per neuron, scrambling the 0’s and 1’s across all time, therefore preserving the global firing rate of a neuron. The second method shuffled the binarized data across the recorded neurons in a given instant of time, keeping the number of spikes at a given time the same, but changing identity of the spiking neuron. Both methods were done and are labeled collectively as “Shuffled Spikes” throughout. State tables were also shuffled, as in Figure 3, to produce the random set for SSI. This shuffling was similar to method one for spiking, wherein substate label was shuffled in time for a given feature. Both of these shuffling methods also produced the shuffled dataset for hub computation (see Fig. 5).

State classifier. To understand how many neurons should be simultaneously measured to reliably estimate the current firing, storage, or sharing state, we constructed a machine learning classifier predicting the label of the current state based on features from different numbers of

single neurons. We then tested how the prediction performance varied with an increasing number of included neurons. Specifically, we trained a standard k -Nearest Neighbors classifier (for $k = 3$) using as input feature values from 1 to 50 neurons. To make an average and robust representation of subsamples of neurons, dozens of random combinations were independently run in the classifier for each number of neurons used. The final prediction accuracy for a specific number of neurons was taken as the average accuracy of all computed combinations (for further details on how these combinations have been made, see below). The variance in accuracy was also computed across each random combination and represented by error bars in the graphs. For each combination, the classifier was trained on a fivefold validation process: meaning that 80% of time points were used for training and accuracy of prediction was tested on the resulting 20%. This was done 5 times, changing the train and test time points every time. Additionally, data were stratified proportionally between train and test sets according to the occurrence of substates.

To create the combinations of neurons, we used the following process. When increasing the number of neurons, the variance between random combinations quickly drops. Thus, less random combinations are needed when predicting with high number of neurons. For prediction with 1–11 neurons, we took a number of random combinations of neurons equal to half of the total of neurons in the recording. For example, if trying to predict with 4 neurons in a file containing 100 neurons, we used 50 random combinations of 4 neurons among the 100. Importantly, in every combination, one different neuron is fixed to have at least one different neuron present in at least one batch. For prediction with 12–19 neurons, the number of random combinations is half of the previous case. For example, if there were 100 neurons in the file, there would be only 25 random combinations. For prediction with 20–39 neurons, the number of random combinations is divided by 3 compared with the first case of 1–11 neurons. For example, if there were 100 neurons in the file, there would be only 16 random combinations. Finally, for prediction with 40–50 neurons, the number of random combinations is divided by 4 compared with previous case with 1–12 neurons. For example, if there were 100 neurons in the file, there would be only 12 random combinations.

Contrast. To calculate contrast for a given feature, we first calculate its global mean for each neuron (i.e., global mean firing per neuron). Here, “global” refers to the entire recording. We then calculate the substate mean for each neuron by concatenating all periods of a given substate and calculating the mean across the “entire” substate. The formula for contrast is then defined as the difference between the substate mean firing rate and the global mean firing rate, normalized by the global mean firing rate as follows:

$$\text{contrast} = \frac{\mu_{\text{substate}} - \mu_{\text{global}}}{\mu_{\text{global}}}$$

This allows the contrast to be either positive or negative. This process was done for all three features of firing, storage, and sharing such that there are contrast values for each. This process was repeated for all possible clustering, therefore a contrast value per feature per k .

Substate specificity. To compute the distribution of substates within periods of SO and THE, we counted the number of times a substate appeared within a given epoch. Some substates exclusively appeared in only SO or THE, while others occurred in both. From these frequencies, we estimated $p(\text{THE})$ and $p(\text{SO})$, that is, the probability of a given substate occurring in either THE or SO, respectively. SSI is then:

$$\text{SSI} = |p(\text{THE}) - p(\text{SO})|$$

This equation results in SSI bound between 0 and 1, where 1 represents a state who exclusively occurs in either THE or SO and 0 represents a state that occurs equally in THE and SO.

Hubs and hub stability. In this work, we define a hub neuron in the same way as our previous work. Namely, for a given feature, if a neuron’s activity within a given substate was higher than the 90th percentile, it

was marked as a hub for the feature for that state. We compute hubs for every iteration of state table as defined above, such that we have a graph, or matrix (see Fig. 4A) for each state table. These matrices are Neuron \times k_{total} where each entry is either a 0 for non-hub or 1 for hub. To compute how stable each of these matrices are as a function of k , we compute the normalized hamming distance of each matrix using the *pdist2* function in MATLAB but modified so that it gives a sense of how stable hubs are across states, where perfect similarity would result in a 1, and no similarity at all would give a 0.

Dictionary and complexity. To compare sequences of substates of different types or in different regions, we introduced a symbolic description of substate switching. With this description, each substate label acts as a letter symbol $s^{(p)}$, where (p) can indicate firing, sharing, or storage. For example, the firing features from the example substate table 1 (see Fig. 2A) would have the integer labels 1, 2, 3, and 4 (they can also arbitrarily be assigned letters as well, i.e., A, B, C, and D). We can therefore describe the temporal sequences of the visited substates of each feature as an ordered list of integers $s^{(p)}(t)$. Once substate labels are thought of as letters, we define the combination of firing, storage, and sharing letters in each state table from a given window as 3-letter words. Using the formalism of linguistics, we can then compute the dictionary, or the number of words expressed, of a given recording within a region. We can also compute the used dictionary fraction, or the number of words found in the dictionary divided by the number of theoretically possible words given the number of substates per feature. For example, substate table 1 could have expressed 27 unique words. The used dictionary fraction was computed in an identical way to that of Clawson et al., 2019 (see Complexity of substates sequences).

Using these methods, we compute the complexity of the sequences expressed using the notions of Kolmogorov–Chaitin complexity and minimum description length approaches (Crutchfield, 2012). While further discussion of method can be found here (Clawson et al., 2019) — the aspects of this complexity measure that is relevant for this work is that a random sequence of letters (and words) produces a higher complexity, while an ordered sequence of letters (and words) would produce a low complexity.

Ordered and random substate tables. To have relevant points of reference in our measures, each substate table was ordered and randomized. For the case of ordering, all substate labels for all features were sorted in ascending order which keeps the total lifetime of any state constant, while removing the temporal organization in an ordered fashion. In the case of randomization, all substate labels for all features were randomized 500 times, which again keeps the total lifetime of any state constant, while removing the temporal organization in a random fashion.

To compute the relative minimums and maximums for comparisons between order and random, the MATLAB function “rescale” was used. The minimums were computed using the average (of a given measure) of all ordered state tables for a given k_{total} , and the maximums were computed using the average (of a given measure) of all random substate tables for a given k_{total} .

Plotting. Various tools were used for plotting with most done via MATLAB, other tools were also used from “Moving Beyond p values” (Ho et al., 2019).

Data availability. Partial data and codes can be found here: 10.5281/zenodo.4534369. Full codes, including figure generation as well as complete dataset, are available upon request.

Results

Alterations of firing dynamics in THE and SO

We sought to first establish differences in neural activity with respect to periods of THE and SO for both control and epileptic animals. To do this, we analyzed LFPs and action potentials from individual neurons measured from the HPC (CA1) and mEC from control ($n = 5$) and experimental epilepsy ($n = 6$) rats under anesthesia (Fig. 2A; for details, see Materials and Methods). We chose to assess basal network dynamics under

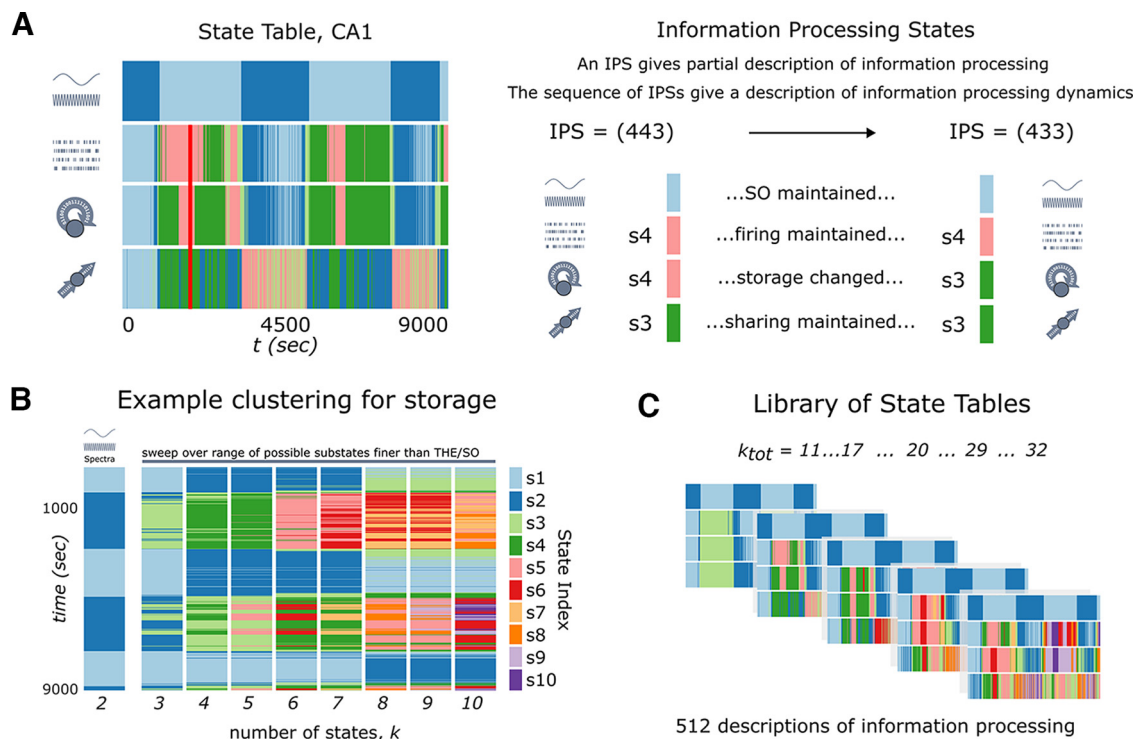


Figure 4. Information processing states and state tables. **A**, An example of a state table for CA1 from a control animal. State tables are symbolized representations of ongoing information processing: each feature captures a separate aspect of the information processing and a vertical “slice” of this state table is an IPS. We suggest that IPSs capture a partial description of information processing and their sequence is therefore a partial description of their dynamics. Interestingly, substates do not all change simultaneously, as in the example on the right. At a given time, t , IPSs are labeled by a 3-letter word, each letter corresponding to the substate label for firing, storage, and sharing. **B**, An example of the clustering method used, wherein a range of k values are selected. As k is increased, the storage feature is clustered ever more finely, introducing a new substate at each step, and substates are color-coded as presented in the legend. Moving from left to right the finer scaled clustering can be seen by eye. While colors may change from one k value to the next, the structure of the clustering is what is key. **C**, An example of a state table library, or the entire collection of possible information processing descriptions possible using our methodology.

anesthesia and not during sleep to remove a potential confounding effect of seizures and interictal events that occur during sleep in this model of epilepsy (see Materials and Methods). While both regions were recorded simultaneously, they are analyzed separately, as we focus on local information processing within regions. Unsupervised clustering of the spectral content of LFPs reveals that field activity continuously switches between two states: SOs (0.5–3 Hz) and THE oscillations (3–6 Hz) (Fig. 2A). As previously reported in freely moving animals (Chauvière et al., 2009), THE power and peak frequency are decreased in CA1 in experimental epilepsy. While the peak frequencies of THE and SO are not modified in the mEC in epilepsy, their power is decreased. However, the ratios of frequencies and powers across SO and THE are similar in both control and epilepsy, suggesting a conservation of underlying processes (Extended Data Fig. 2-2).

Next, we examined whether firing, storage, and sharing (Fig. 1) are dependent on the brain state (THE vs SO), the region (CA1 vs mEC), and the condition (control vs epilepsy). In control animals, we find a higher average firing during THE than SO in both regions. However, we see no significant difference in information processing metrics between THE and SO, except for storage in CA1 (Extended Data Fig. 2-2). In epileptic animals, there is a decrease (with respect to control) in all spiking features during both THE and SO epochs and in both regions (Extended Data Figs. 2-4, 2-5).

Specifically, firing rates in epilepsy were slightly reduced; however, the effect was small, with a general average reduction of $\sim <1\%$ (not significant for both regions). More marked was the reduction in the difference between firing in THE and SO

epochs, as this difference almost disappeared in TLE. Storage and sharing held the same relationship as in control, except in the case of storage in CA1 (for p values, see Extended Data Fig. 2-6). Overall, the maintained ratios for frequency, power, and information processing metrics across THE/SO epochs suggest that, while firing is altered in experimental epilepsy, there remain properties or rules that are retained despite this alteration. In other words, the inspection of average properties does not reveal striking qualitative differences between the control and epileptic conditions. These analyses, however, focus on mean values across global brain states. Differences between control and TLE conditions may arise at the subtler level of how primitive computations are organized in space and time within each of these global states as we address in the following.

Control and epileptic activity can be described in terms of dynamically switching substates

Following Clawson et al. (2019), we identify substates for firing, storage, and sharing features. The substate detection method is based on finding periods of highly correlated feature values across the recorded population of neurons, computed using neuronal spikes (for number of neurons recorded, see Materials and Methods). Figure 2 and the following text gives an example for firing substates. The process of substate extraction is identical for the storage and sharing features (Fig. 2B; for details, see Materials and Methods). The firing rates of all neurons are calculated in a 10 s window, the window is shifted 1 s, and the rates recalculated. This process is repeated for the entire recording, resulting in a data matrix that is neurons \times windows in size. Then, a similarity matrix is created by calculating the

correlation between all pairs of windows in a given recording, resulting in one of the two graphs in Figure 2C (i.e., each graph is computed separately; see Materials and Methods). A blocky structure is evident in both graphs by eye, with bright yellow blocks representing periods of time wherein the firing dynamics across the recorded population within a region is similar. The existence of these blocks was not granted (similarity matrices could have been unstructured) and proves thus that substates are not uniquely a mathematical abstraction but correspond to a distinctive aspect of firing organization. This goes beyond mere changes of average firing rate, as population firing rates may be constant, but which neurons are spiking can change. In this example, the yellow blocks are closely aligned to epochs of THE and SO because of the relationship between firing rate and global oscillation; however, as highlighted in the red square, it is often the case that there are “blocks within blocks” wherein the activity switches between two different modes of firing despite the global state (oscillatory epoch) staying the same. We then run an unsupervised clustering algorithm, k -means, to extract these blocks, with the results shown below each graph for an example k . Overall, Figure 2 establishes that feature values (here firing) can remain stable across the recorded population during a given time period (i.e., during successive windows), before switching to a different set of feature values for a new period of stability. Here, we identify (and later analyze) substates of firing, storage, and sharing with the clusters found by k -means. Importantly, k is a parameter that can be changed and will determine how many substates are found; and in Figure 2C, we show 5 substates in control (left) and 4 in epilepsy (right). In later figures, we choose k values that are easy to visually represent to illustrate the measurements we take that rely on substate labels. However, we test many values of k , and therefore the number of resultant substates, in all measurements (see Extended Data Fig. 3-1). We discuss this choice in the following section.

The detection process done for the firing feature can be repeated as well for the other storage and sharing features, finding a qualitatively analogous block organization into similarity matrices and, correspondingly, the existence of switching between substate epochs. Firing substates, as discussed above, represent periods in which specific neuronal subsets are consistently firing at a higher firing level. Analogously, storage substates represent periods in which specific neuronal subsets display stronger involvement into information storage. It is key to note that substates of storage can switch despite the firing substate remaining stable. This would indicate that while the same neurons may be active, potentially even with the same mean firing rate, the temporal structure of their firing (e.g., interspike interval) may change, as reported in Clawson et al. (2019). Sharing is altogether different, as it is a network measure, rather than an individual neuronal measurement, as in firing and storage (see Materials and Methods). Sharing substates therefore represent periods wherein sharing networks (which neurons are sharing with whom and how much) are highly correlated. Further confirming Clawson et al. (2019), we find that substates of sharing can switch despite the firing and storage substates remaining stable. While this organization of firing, storage, and sharing through time is the focus of this paper, we first investigate neuronal activity inside of these substates, in both control and epilepsy.

Information substates are more contrasted in epilepsy

To quantify the substates internal “shape,” we developed a measurement of contrast. Taking storage as an example, we first calculate a neuron’s global mean storage (over the whole duration of the recording), and its mean storage within each substate. The

contrast is defined as the difference between the substate mean storage value and the global storage value, normalized by the global storage value, for each neuron (Fig. 3A). Contrast thus gives a measure of how much a given feature (here storage) is changing for a neuron within a given substate, with respect to average levels. This measure of contrast can be applied across all features, and Figure 3B shows the contrast plots for an example recording, containing 55 recorded neurons in CA1. The number of substates in this example were chosen for easy visualization, as the number of states can be adjusted within the method and is discussed in the next section. The differences between substates now clearly appear as large changes in the profiles of contrast values for the example recorded neurons. To quantify this, we computed the mean contrast of each substate for each feature, that is, the average of the absolute values of the heights of the bars in the contrast plot (data not shown, see Materials and Methods). Substates with high average contrast are more distinctly detaching from the background of globally averaged behavior. Figure 3C shows the distributions of the bootstrapped mean differences of contrast between control and epilepsy for firing, storage, and sharing features in CA1 using the same substates as Figure 3B (see Materials and Methods). While contrast is higher for firing in control, both storage and sharing have increased contrast in epilepsy. This suggests that, while neuronal firing in epilepsy generally deviates less from the global mean than control, substate switching more strongly modulates storage and sharing in epilepsy. We thus identify one major alteration in epilepsy at this algorithmic level; information-theoretic substates are more contrasted, exhibiting more marked differences with respect to the mean, than their control counterparts. Substates are thus not suppressed in epilepsy with respect to control. On the contrary, they are even more markedly distinct. When k is varied, this relationship is held across the tested range for all features in both regions, although less drastic in mEC (Extended Data Fig. 3-1).

Loss of global state specificity of firing and storage substates in epilepsy

We then determined the brain state specificity of the substates for firing, storage, and sharing features in control and epileptic conditions. We calculate the state specificity index (SSI), the probability that a substate occurs during a THE epoch, SO epoch, or both. The SSI is a metric bounded between 0 (a substate occurs equally in THE or SO) and 1 (a substate is exclusive to either THE or SO) (see Materials and Methods). Figure 3D shows an example of this measure and the distribution of mean SSI for substates in control (blue) and epilepsy (red) recordings in the case of $k = 4$ for firing and sharing and $k = 3$ for storage (as in Fig. 3B). Firing and storage have SSI near 0.8 in control animals, indicating high brain-state specificity, whereas the SSIs are much lower in epileptic animals. Sharing shows no significant shift in SSI between control and epilepsy conditions. All values are well above the null models created from shuffling (gray) and sorting the tables (black) (see Materials and Methods), which indicates that, even in epilepsy, the nesting relationships between global states and substates, while diminished, are not totally unstructured. This trend holds across both mEC and CA1, as well as across a range of k , or number of substates (Extended Data Fig. 3-2). We interpret this finding as an indication that the substate hierarchical articulation to brain states becomes more “disordered”; that is, a large proportion of firing and storage substates now occur during both THE and SO in both mEC and CA1 in epilepsy. We did not find any reduction of SSI in epilepsy

for sharing substates. However, for sharing, SSI values were low already in the control case, and the absence of epilepsy-related modulations may thus be because of a saturating “floor effect.”

Together, our results show that the involvement of different neurons in information processing is more heavily affected by substate switching (contrast increase) and the substates themselves are no longer as constrained by global oscillatory state (SSI reduction). Thus, epilepsy is characterized by “disorder” compared with control.

Information processing states

Until now, results have been presented on a feature-by-feature basis. For a more complete view of information and its dynamics, all features need to be considered simultaneously. We therefore revisit here the concept of an information processing state (IPS), that is, a combination of temporally coinciding substates of firing, storage, and sharing in a recording (Clawson et al., 2019). The notion of IPSs allows high-level analyses of algorithmic organization without need to refer to actual values of information processing features. The description is “compressed” into just the sequence of the visited IPSs and their identity at any time, akin to symbolization approaches in nonlinear physics (Porta et al., 2015).

Figure 4A illustrates IPSs with an example of the collection of substates of the three features alongside the global state (the state table) and an IPS (red vertical line) in a ~2-h-long recording performed in the CA1 area in a control animal. From each analysis time window, we generate a three-letter word, with the letters representing the substate labels of firing, storage, and sharing features measured in this time window (see Materials and Methods). When the analysis window is shifted by 1 s, another word is obtained, which is identical to the previous one if the labels (and therefore substates or global state) do not change. This procedure allows us to reduce the description of the complex simultaneous variations of firing, storage, and sharing patterns within the neuronal population to simple strings of symbolic words. We can then assess how the properties of these strings are modified in epilepsy. Therefore, we present IPSs as a partial description of primitive information processing during a window of time, and the sequence of IPSs through time in a state table as a description of information processing dynamics.

As substates represent stable periods of neuronal activity and are the core of the IPS, the number of substates for a feature (i.e., the k value) in any given recording is a free parameter, which determines the number of IPSs. The ground truth, how many clusters exist, is unknown. Statistical criteria were previously used to find their (supposedly) optimal number (Clawson et al., 2019). Here, we use a generalized and unbiased approach, varying the k value for each firing, storage, and sharing feature while fixing $k = 2$ for the spectral feature and study the dependency on k of the different metrics of interest. Figure 4B illustrates this concept, showing the estimated storage substates through time as k increases from 3 to 10. A low value may potentially underestimate the real number of substates, while a large number may be an overestimate producing substates that rarely occur more than once (see Materials and Methods). We therefore use a lower bound of $k = 3$, and a reasonable upper bound of $k = 10$, wherein the clusters become too fine (see Materials and Methods). Each triplet of k values, regardless of the k used, will produce a state table that represents a specific symbolization of a recording (Porta et al., 2015). These tables can be used not only to examine the neural activity within the states, as in contrast, but also the structure of the state transitions, which is discussed further in the paper. In both control and epilepsy, the silhouette difference

between the clusters, a common measure of “goodness of cluster,” remains well above a null hypothesis silhouette difference for all features in both regions (see Materials and Methods; Extended Data Fig. 3-3). To understand the effect of number of neurons recorded and substate estimation, we built a classifier (3-nearest neighbors, see Materials and Methods) that used a subset of the recorded neurons’ feature values to predict substate label. We observe that for all features ~20 or more neurons were enough to accurately predict substate label (Extended Data Fig. 3-4). This is somewhat in agreement with previous work detailing network states based on pairwise measurements which predicted 10 needed neurons (Schneidman et al., 2006), and this result held in both regions across control and epilepsy.

As we are interested in understanding IPSs and their sequences, the three features are considered simultaneously, and each triplet of k values will produce a specific set of IPSs and state table. Together, we consider eight possible k values for each feature, giving rise to $8^3 = 512$ possible state tables. Each state table is characterized by the total number of substates it contains: $k_{tot} = 2 + k_{firing} + k_{storage} + k_{sharing}$ with a maximum value of $k_{max} = 32$ ($32 = 2$, the number of spectral states + 3 features \times 10). The collection of all state tables for a given recording defines a library of tables (Fig. 4C). We chose such a method with the intention that without an *a priori* approach on the underlying principle, if we extract generic rules, they should be valid independently of the choice of number of clusters, at least for a reasonable wide range of k values. All analyses that can be done on a state table are performed for each library, across the range of k_{tot} , which gives an added benefit of assessing the robustness of the results regardless of the number of clusters.

Computing hubs are more numerous but less substate-specific in the mEC in epilepsy

To further emphasize unique contributions of features to the IPS, we extract computing hub neurons. These are neurons with on average, exceptionally high firing, storage, or sharing values with regard to a given substate (see Materials and Methods). The logic here is to see whether a small set of neurons are “responsible” for substate behavior or substate switching. Different substates can be associated to different sets of hubs (Clawson et al., 2019). Neurons acting as a firing, storage, or sharing hub in a given substate will not necessarily do so in another substate. Therefore, while the fraction of neurons being hub in a given substate remains small, the fraction of neurons serving as hub at least in one substate is much larger, approaching ~40% on average. Figure 5A illustrates an example of the distribution of hubs (same recording as in Fig. 4A, with $k_{tot} = 14$).

We observe that the percentage of neurons serving as hubs at least once is increased in both regions in epilepsy (this increase, however, is statistically significant only in CA1; Fig. 5B). This result is well compatible with the increase in substate contrast found in epilepsy for all k values (Extended Data Fig. 3-1): as more neurons exhibit a large contrast, more of them can also be detected as hubs. For both control and epilepsy, the percentage of neurons marked as hubs is significantly larger compared with randomized state tables (Fig. 5B, gray distribution).

Figure 5A also shows that some computing hubs are shared by different substates, while others are specific to one substate/one feature. In order to assess how substate-specific the computing hubs are, we use a measure of similarity (see Materials and Methods). Here, 0 indicates that every substate has a unique hub set with no overlap between substates, while a 1 value means that

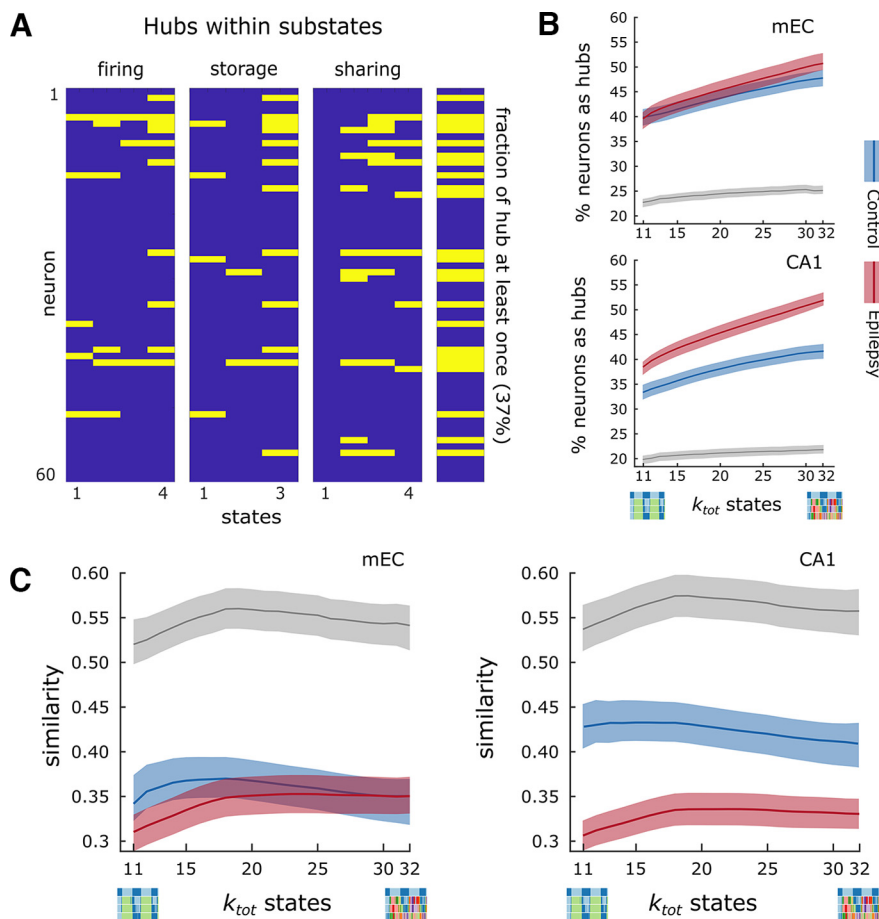


Figure 5. Computing hubs and their distributions. **A**, Example of computing hubs in control CA1 extracted from a $k = 4$ (for the three features) state table. y axis indicates unsorted neuron label. x axis indicates the substates for firing, storage, and sharing features. Yellow bar represents that the given neuron is a computational hub during a substate. Right, Summed version of the graph on the left, visually showing the fraction of neurons that are a hub at least once (37%). **B**, The percentage of neurons that are a hub at least once is increased in epilepsy independently of k_{tot} . Gray distribution represents the mean of the shuffled, null model. **C**, The similarity index plotted as a function of k_{tot} . The hubs become more substate-specific in the CA1 in epilepsy. Blue and red represent control and epilepsy data, respectively. Bold lines indicate the mean. Shaded regions represent the 99% bootstrapped CI. Gray distribution represents the mean of the shuffled, null model. The similarity index is plotted as a function of k_{tot} .

all substates have an identical distribution. Figure 5C shows that, in control animals, a majority of hubs tend to be substate-specific (similarity < 0.5). In CA1, the distribution of hubs is less substate-specific than in the mEC (higher similarity). In epilepsy, the distribution of hubs does not change in mEC, while hubs become significantly more substate-specific in the CA1. In other words, the status of being hub is for a CA1 neuron less stable in epilepsy than in control animals.

We conclude that, in epilepsy, the mEC and CA1 display an increase in the number of neurons labeled as hubs at least once, and that the substate specificity of hubs is increased in the CA1. In mEC, there are only fractionally more hubs and they are similarly distributed across substates as in control. In CA1, there are many more hubs, and these hubs are more “unique,” as they overlap significantly less across substates as in control. These hubness analyses provide further indications of increased “disorder” in epilepsy, as information processing lead roles are more chaotically distributed across available neurons.

Increased dictionary in epilepsy

We described state tables as sequences of words and all possible state tables generated through our k -means procedure as a library (Fig. 4). While we do not commit on the fact that IPs and state tables are genuine ingredients of an actually used

“language of neurons,” we can still use this linguistic metaphor to inspire expansions of our analytic toolbox. For example, if each state table is seen as a book, we can explore the structure of IPs sequences, examining the vocabulary extent and style of all the “books within the library.” What words are expressed? Is there a syntax, or organizational rules? And how does epilepsy change these measures? To this end for each library, we build two sister libraries for comparison: one in which we sort every book internally to be highly ordered, and one in which we randomize every book internally to be highly disordered (Fig. 6A; see Materials and Methods).

For each k_{tot} , there is a fixed number of potential words that can be generated and possibly appear within the associated book. As in any language, only a fraction of all possible words is expressed. For each book, we measure the used dictionary fraction, or relative dictionary (see Materials and Methods). Figure 6B illustrates two end cases with a diagram. The low relative dictionary case (left) uses a small number of expressed words, while the high relative dictionary case (right) uses a much richer vocabulary, wherein almost all of the potential dictionary is expressed. While the measure of relative dictionary in and of itself is informative, it is difficult to use such a measure to assess meaningful changes (i.e., between control and epilepsy) without having comparative baselines. Therefore, we compute not only

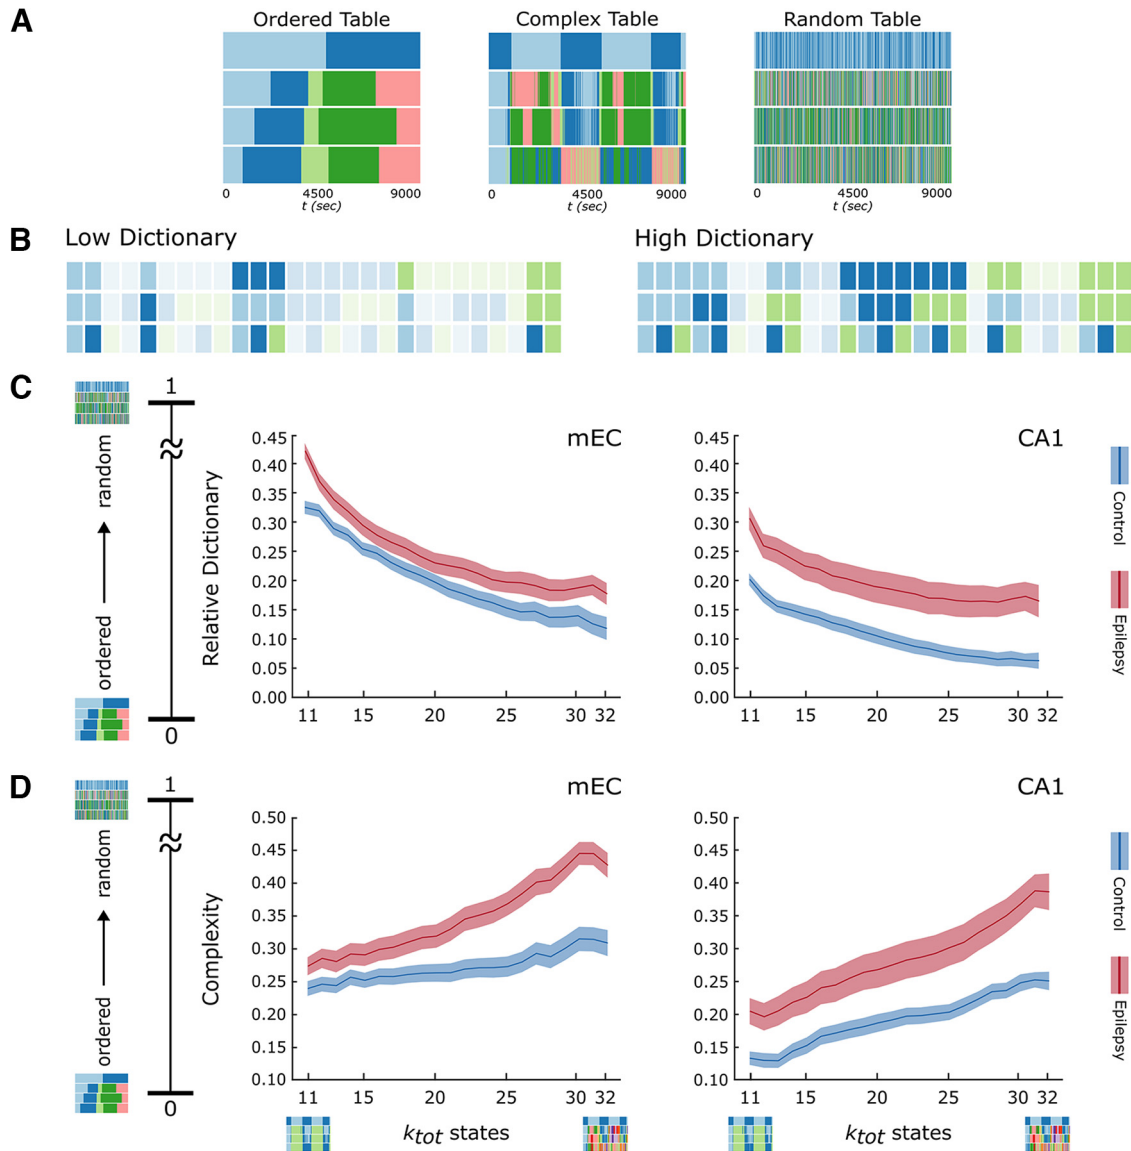


Figure 6. Dictionary and complexity in control and epilepsy. **A**, Examples of state tables, constructed from data from CA1, showing the two extremes of order and disorder as well as one of the possible state tables taken from the state table library. **B**, Fictional diagrams represent two extremes for the measure of relative dictionary. Each row represents a feature (firing, storage, sharing); for simplicity, we do not consider the brain states (THE and SO). We consider three substates (light blue, dark blue, green) per feature (using the same color code for simplicity), which makes a total of $3^3 = 27$ words (the representation is similar to counting in base 3 with color, increasing from left to right). Words that are not observed are shaded. A low relative dictionary (left) contains a low fraction of all possible words, while a high relative dictionary (right) contains a high fraction. **C**, Relative dictionary values as a function of k_{tot} . As expected, the fraction of words used in control decreases as the number of possible words increases. The relative dictionaries are similar in mEC and CA1 in controls. There is a marked increase in the relative dictionary in both regions in epilepsy. Blue represents control data. Red represents epileptic data. Bold lines indicate the mean. Shaded regions represent the 99% bootstrapped CI. **D**, Complexity values for both the mEC and CA1 as a function of k_{tot} . The complexity is similar in mEC and CA1 in controls. In epilepsy, the complexity is increased in both regions.

the relative dictionary of our libraries, but also that of the ordered and random paired libraries. These correspond to the null hypotheses of order and disorder in the “language” of the book, respectively. Intuitively, ordered books have few words that are repeated in large blocks of time, whereas disordered books have many words that are repeated by chance. We then apply a linear transformation to the relative dictionary measure, such that a rescaled value of 0 represents the relative dictionary measure expected on average for ordered books, and a rescaled value of 1 representing a relative dictionary measure identical to that of randomized books. Such a normalized relative dictionary measure tracks not only the richness of the used dictionary but also its positioning between order and disorder.

Figure 6C shows that, for both the mEC and CA1 in control and epilepsy conditions, the normalized relative dictionaries

lie much closer to 0 than to 1, indicating that, in all cases, their relative dictionaries are closer to systems with an ordered rather than disordered organization. In epilepsy, the relative dictionary is increased with respect to control in both mEC and CA1 (Fig. 6C).

This increased dictionary is not because of epilepsy having more states than control. This measure captures that, given the possibility to express an equivalent dictionary size (by fixing k_{tot}), epileptic recordings always express a higher fraction of this dictionary than control. This additionally supports the hypothesis of increased disorder, with one potential hypothesis being that the stability of states in control is decreased in epilepsy, allowing the neurons to exhibit a wider variety of IPs, many of which could be transient “noise glitches” rather than proper functional configurations (see Discussion).

The relative dictionary provides important information about which words are found, but not how words are organized in time. This is like the grammar, or syntax, of a traditional sentence and to analyze this syntax (i.e., the rules according to which words concatenate from one time window to the next), we quantify the level of organization present in the state tables as a whole (i.e., the overall dynamics of a system moving through IPSs) (Fig. 4).

The syntax of substate sequences is less regular in epilepsy

Compressibility is a key property of a (symbolized) object as it represents its degree of internal order. This is because any regularity within the object may be described by simply referencing to its previous occurrences. As in relative dictionary, we consider state tables bordered by two extreme cases: order and randomness (Fig. 6A). An ordered table is dominated by a highly structured syntax, typically dominated by a lower dictionary and long periods of sustained words. Therefore, an ordered table is very compressible because of this internal regularity. A random table, on the other hand, typically contains an exceedingly high number of words, which follow each other in a disorderly (random) manner. This results in noncompressibility. A complex table is one that lies between those extremes. In order to characterize the complexity of the state tables, we compute the Lempel-Ziv complexity of the state table (Rissanen, 1978; Clawson et al., 2019), which is scaled to the sister libraries of order and disorder. Thus, in Figure 6D, a rescaled table complexity of 0 represents the complexity of books within the ordered library, something very compressible, while 1 represents the expected complexity for randomized books in our disordered library, something very uncompressible (as shown in Fig. 6A). In controls, the complexity is higher in mEC than in CA1 and closer to ordered than to disordered tables (see Discussion), as for the relative dictionary. In epilepsy, the complexity is significantly increased for both regions across all k_{tot} values.

We therefore propose the following interpretation: neuronal activity, as characterized through action potentials, is altered in epilepsy in a way that not only produces more varied substate behavior but alters the way in which the substates combine. This drives the neuronal system into a regimen where organizational relationships are less stable, allowing for a wider but less coordinated repertoire of firing patterns (and their associated algorithmic effects), perhaps detrimental to system's function as a whole.

Discussion

This study provides evidence that epileptic conditions alter information processing in its simplest sense, the primitive storage and sharing operations as we introduce here, in both the mEC and CA1. As these basic processes may underlie a variety of neural computations, their alterations may impact numerous cognitive functions. Looking at simpler, classical properties, the differences between control and epileptic conditions were not striking. Power and firing rates were slightly reduced in epilepsy (Extended Data Figs. 2-2, 2-3); however, the effects were small and, in the case of power, did not affect the relative ratio between THE and SO bands. It was only when looking at subtler aspects of firing organization, such as its articulation into a multiplicity of states forming complex sequences, that differences between the two conditions became more drastic. Our findings thus suggest that epilepsy may be affecting the “informational effects” “of collective neuronal activity more

than” activity itself, at least in interictal periods far from seizure events.

The main limitation to our study is that it is conducted under anesthesia versus, for example, goal-directed behavior to assess cognitive function. The type of analysis we performed is powerful as it allows unraveling basic properties of information processing without needing to know which computations are ongoing. However, it requires long-duration, stable recordings with large state sampling to obtain enough data points to perform reliable statistics. We did not record during natural sleep, as seizures and interictal spikes (which would act as strong confounding factors) mostly occur during the light phase, while they do not occur under anesthesia. However, a similar type of analysis performed in control animals led to similar results during sleep and anesthesia (Clawson et al., 2019), suggesting that the anesthesia procedure we use does not significantly alter core information dynamics (at least with respect to natural sleep, if not awake behavior, in control animals).

We refer to the elementary information storage and sharing operations as primitive (or low level) information processing operations, as we consider them as fundamental building blocks within an algorithm to reach an end condition (like a function), similar to the “algorithmic level,” introduced by Marr and Poggio (1977). Algorithm is used here in its most generic meaning, as we do not claim that the brain is analogous to an electronic computer. Such primitive processing operations, as we define them, represent nothing else than the emergent “informational effect” of very concrete neurophysiological phenomena. Storage and sharing of information directly derive from auto- and cross-correlations in firing, which widely vary in neuronal populations (Schneidman et al., 2006), and can be directly measured from spiking activity. Other primitive processing operations exist, such as information transfer (Schreiber, 2000; Palmigiano et al., 2017) or information modification (Lizier et al., 2013; Wibral et al., 2017). Our recordings and choice of a time-resolved approach do not provide enough data to track these more sophisticated operations. It is for this reason we refer to IPSs as a partial description of information processing (Fig. 4). However, the processing functions of storage and sharing are important as they represent statistical measures of information maintenance and spread in time and space time, respectively.

We show that primitive information processes are organized in temporal sequences of IPSs, which are extracted via a cluster analysis. We have used a nonbiased approach, spanning many possible combinations of numbers of clusters. The fact that most results are independent from the choice of the number of clusters provides a strong argument for the genericity of our conclusions. With this approach, we demonstrate a degradation of algorithmic organization because of enhanced randomness in epilepsy. This conclusion stems from the convergence of complementary analyses.

First, the state specificity of IPSs is reduced; that is, many IPSs are now redundant between THE or SO. Second, the contrast of storage and sharing states is made more erratic in epilepsy, fluctuating around their means much more than in control (and this despite firing rates being, on the contrary, less modulated; compare Fig. 3), suggesting subtle alterations of inter-neuron communication. Third, the increasing numbers of storage and sharing hubs, especially in CA1, paired with the decrease of state specificity, demonstrate that substates are more different from one another than in control and less constrained by brain state. Fourth, freed from the constraint of being strongly state-specific,

the relative dictionary in epilepsy is increased. Finally, the description complexity of IPS sequences tends to be larger in epilepsy than control. In other words, IPS sequences have a less regular syntax, potentially stemming from the increased repertoire of IPS states. These results are also compatible with the degradation of dynamic rich club organization of information sharing previously observed in TLE relative to control (Pedreschi et al., 2022).

The fact that IPS dynamics across mEC and CA1 are different in control and that information processing is affected in brain region-dependent manner is an important result. The IPS dynamics of CA1 show, in general, more alterations than that of mEC. The mEC and CA1 have distinct cytoarchitectures and different fates following an epileptogenic insult. Most striking is the loss of layer 3 in the mEC, and the injury of many pyramidal cells and interneurons in the CA1 region (Curia et al., 2008). At present, it is not possible to assign a given alteration in information processing to particular anatomic-functional changes in the mEC or CA1.

Global brain states (THE and SO) and IPSs are emergent properties. Any change in any brain region can potentially affect neuronal dynamics anywhere from the local to the global scale (Rocchi et al., 2022). Therefore, the anatomic-functional alterations in mEC or CA1 may contribute to any combination of local and global changes. Changes in terms of information processing do not necessarily have to be homogeneous across brain regions. Indeed, brain region-specific modifications are expected as each region is embedded in different functional networks. How these brain region-specific changes contribute to comorbidities (e.g., cognitive deficit, anxiety, and depression) remain to be determined.

Our measure of complexity is that of compressibility, accounting for the internal structure (i.e., how internally ordered are IPS syntaxes). Any change in this internal organization would thus imply an underlying change in the capability for algorithmic operation, possibly resulting in different computation in control and epilepsy conditions. Our measure of complexity does not allow distinguishing between an increase in processing versus an increase in noise, as complexity would grow in both cases. Other measures can be used, but they would require more data (Crutchfield, 2012). Therefore, it is impossible to conclude whether this increase in complexity is “negative” (more disordered information processing) or “positive” (more elaborate information processing). Indeed, some studies have shown enhanced complexity in states with enhanced computation: complexity is boosted both in REM (bearing similarity with THE) versus non-REM (SO) sleep (Clawson et al., 2019) or in conscious versus unconscious conditions (Wenzel et al., 2019). We cannot exclude that enhanced computation in TLE is a positive adaptation to compensate for other deficits associated with the pathology. However, we find more likely that, in epilepsy, the increase in relative dictionary and the drift of complexity toward randomness may reflect a more irregular and chaos-like ergodic IPS selection, with rare IPSs indicating dysfunction in IPS sequential production. At face value, these results may seem intuitive — injury and damage result in less constrained dynamics in the same way a damaged car would drive worse. However, epilepsy is a disease commonly marked by hypersynchronous activity, which could indicate even more constraints than pre-injury. This work’s goal is to demonstrate that, while synchronous activity may be common, global states such as THE/SO epochs preserve in epilepsy rich repertoires of collective spiking activity patterns, very similar

to control conditions. However, subtle modifications in this repertoire and in the way into which the system samples it jointly build up into profound alterations of what we call primitive information processing.

In a biological context, such algorithmic level changes could lead to an entirely different expression of higher-level behavior, such as cognition. However, the question of whether the increase of complexity (decrease of internal order) observed in epilepsy is the source of cognitive deficits or not remains open. It has been theorized that “biological systems manipulate spatial and temporal structure to produce order – low variance – at local scales” in an effort to adapt and survive (Flack, 2019). Therefore, if networks are still functional in epilepsy conditions, are these manipulations now less effective? Or is the resulting low variance order now too difficult to sustain because of a combination of physiological and functional changes? These issues remain to be addressed. Nevertheless, the approaches presented here introduce valuable insight into aspects of the collective behavior of neural populations, and provide a quantitative framework to answer such questions.

In conclusion, the framework we introduce here to compare information processing between control and epilepsy could be generalized in perspective to other neurologic disorders and conditions as well. For instance, aging and Alzheimer’s disease have also been associated to variations in the complexity and orderliness of neuronal activity (although probed at a different macroscale ignoring the finer scale of neuronal firing) (Battaglia et al., 2020; Tait et al., 2020). Following the principle of degeneracy (Prinz et al., 2004), very different structural alterations, which characterize different neurologic disorders, may produce similar alterations in information processing, providing an explanation for the commonalities of comorbidities across different disorders.

References

- Barba C, Barbati G, Minotti L, Hoffmann D, Kahane P (2007) Ictal clinical and scalp-EEG findings differentiating temporal lobe epilepsies from temporal ‘plus’ epilepsies. *Brain* 130:1957–1967.
- Bartolomei F, Chauvel P, Wendling F (2008) Epileptogenicity of brain structures in human temporal lobe epilepsy: a quantified study from intracerebral EEG. *Brain* 131:1818–1830.
- Battaglia D, Boudou T, Hansen EC, Lombardo D, Chettouf S, Daffertshofer A, McIntosh AR, Zimmermann J, Ritter P, Jirsa V (2020) Dynamic functional connectivity between order and randomness and its evolution across the human adult lifespan. *Neuroimage* 222:117156.
- Berchtold NC, Sabbagh MN, Beach TG, Kim RC, Cribbs DH, Cotman CW (2014) Brain gene expression patterns differentiate mild cognitive impairment from normal aged and Alzheimer’s disease. *Neurobiol Aging* 35:1961–1972.
- Blümcke I, et al. (2013) International consensus classification of hippocampal sclerosis in temporal lobe epilepsy: a Task Force report from the ILAE Commission on Diagnostic Methods. *Epilepsia* 54:1315–1329.
- Calhoun VD, Miller R, Pearlson G, Adali T (2014) The chronnectome: time-varying connectivity networks as the next frontier in fMRI data discovery. *Neuron* 84:262–274.
- Chauvière L, Raftai N, Thinus-Blanc C, Bartolomei F, Esclapez M, Bernard C (2009) Early deficits in spatial memory and theta rhythm in experimental temporal lobe epilepsy. *J Neurosci* 29:5402–5410.
- Cho KO, Lybrand ZR, Ito N, Brulet R, Tafacory F, Zhang L, Good L, Ure K, Kernie SG, Birnbaum SG, Scharfman HE, Eisch AJ, Hsieh J (2015) Aberrant hippocampal neurogenesis contributes to epilepsy and associated cognitive decline. *Nat Commun* 6:6606.
- Clawson W, Vicente AF, Ferraris M, Bernard C, Battaglia D, Quilichini PP (2019) Computing hubs in the hippocampus and cortex. *Sci Adv* 5: eaax4843.
- Crutchfield JP (2012) Between order and chaos. *Nat Phys* 8:17–24.

- Csicsvari J, Hirase H, Czurko A, Mamiya A, Buzsáki G (1999) Oscillatory coupling of hippocampal pyramidal cells and interneurons in the behaving rat. *J Neurosci* 19:274–287.
- Curia G, Longo D, Biagini G, Jones RS, Avoli M (2008) The pilocarpine model of temporal lobe epilepsy. *J Neurosci Methods* 172:143–157.
- de Barros Lourenco FH, Marques LH, de Araujo Filho GM (2020) Electroencephalogram alterations associated with psychiatric disorders in temporal lobe epilepsy with mesial sclerosis: a systematic review. *Epilepsy Behav* 108:107100.
- Edelman GM, Gally JA (2001) Degeneracy and complexity in biological systems. *Proc Natl Acad Sci USA* 98:13763–13768.
- Ferraris M, Ghestem A, Vicente AF, Nallet-Khosrofiyan L, Bernard C, Quilichini PP (2018) The nucleus reuniens controls long-range hippocampo-prefrontal gamma synchronization during slow oscillations. *J Neurosci* 38:3026–3038.
- Flack JC (2019) Life's information hierarchy. Santa Fe, NM: Sante Fe Institute.
- Goaillard JM, Marder E (2021) Ion channel degeneracy, variability, and covariation in neuron and circuit resilience. *Annu Rev Neurosci* 44:335–357.
- Gorur-Shandilya S, Cronin EM, Schneider AC, Haddad SA, Rosenbaum P, Bucher D, Nadim F, Marder E (2022) Mapping circuit dynamics during function and dysfunction. *Elife* 11:e76579.
- Harris KD, Henze DA, Csicsvari J, Hirase H, Buzsáki G (2000) Accuracy of tetra electrode separation as determined by simultaneous intracellular and extracellular measurements. *J Neurophysiol* 84:401–414.
- Hazan L, Zugaro M, Buzsáki G (2006) Klusters, NeuroScope, NDManager: a free software suite for neurophysiological data processing and visualization. *J Neurosci Methods* 155:207–216.
- Hesdorffer DC (2016) Comorbidity between neurological illness and psychiatric disorders. *CNS Spectr* 21:230–238.
- Ho J, Tumkaya T, Aryal S, Choi H, Claridge-Chang A (2019) Moving beyond P values: data analysis with estimation graphics. *Nat Methods* 16:565–566.
- Holmes GL (2015) Cognitive impairment in epilepsy: the role of network abnormalities. *Epileptic Disord* 17:101–116.
- Inostroza M, Brotons-Mas JR, Laurent F, Cid E, de la Prida LM (2013) Specific impairment of 'what-where-when' episodic-like memory in experimental models of temporal lobe epilepsy. *J Neurosci* 33:17749–17762.
- Kirst C, Timme M, Battaglia D (2016) Dynamic information routing in complex networks. *Nat Commun* 7:11061.
- Krishnan V (2020) Depression and anxiety in the epilepsies: from bench to bedside. *Curr Neurol Neurosci Rep* 20:41.
- Lee AK, Wilson MA (2002) Memory of sequential experience in the hippocampus during slow wave sleep. *Neuron* 36:1183–1194.
- Lenck-Santini PP, Scott RC (2015) Mechanisms responsible for cognitive impairment in epilepsy. *Cold Spring Harb Perspect Med* 5:a022772.
- Lenck-Santini PP, Holmes GL (2008) Altered phase precession and compression of temporal sequences by place cells in epileptic rats. *J Neurosci* 28:5053–5062.
- Lizier JT (2013) The local information dynamics of distributed computation in complex systems. Berlin: Springer.
- Lizier JT, Atay FM, Jost J (2012) Information storage, loop motifs, and clustered structure in complex networks. *Phys Rev E Stat Nonlin Soft Matter Phys* 86:026110.
- Lizier JT, Flecker B, Williams PL (2013) Towards a synergy-based approach to measuring information modification. 2013 IEEE Symposium on Artificial Life (Alife), 43–51.
- Lopez-Pigozzi D, Laurent F, Brotons-Mas JR, Valderrama M, Valero M, Fernandez-Lamo I, Cid E, Gomez-Dominguez D, Gal B, Menendez de la Prida L (2016) Altered oscillatory dynamics of CA1 parvalbumin basket cells during theta-gamma rhythmicopathies of temporal lobe epilepsy. *eNeuro* 3:ENEURO.0284-16.2016.
- Marder E, Goaillard JM (2006) Variability, compensation and homeostasis in neuron and network function. *Nat Rev Neurosci* 7:563–574.
- Marr DC, Poggio T (1977) From understanding computation to understanding neural circuitry. *Neurosci Res Prog Bull* 15:470–491.
- Mason PH (2015) Degeneracy: demystifying and destigmatizing a core concept in systems biology. *Complexity* 20:12–21.
- Palmigiano A, Geisel T, Wolf F, Battaglia D (2017) Flexible information routing by transient synchrony. *Nat Neurosci* 20:1014–1022.
- Pedreschi N, Battaglia D, Barrat A (2022) The temporal rich club phenomenon. *Nat Phys* 18:931–938.
- Porta A, Baumert M, Cysarz D, Wessel N (2015) Enhancing dynamical signatures of complex systems through symbolic computation. *Philos Trans A Math Phys Eng Sci* 373:20140099.
- Prinz AA, Bucher D, Marder E (2004) Similar network activity from disparate circuit parameters. *Nat Neurosci* 7:1345–1352.
- Quilichini PP, Bernard C (2012) Brain state-dependent neuronal computation. *Front Comput Neurosci* 6:77.
- Quilichini P, Sirota A, Buzsáki G (2010) Intrinsic circuit organization and theta-gamma oscillation dynamics in the entorhinal cortex of the rat. *J Neurosci* 30:11128–11142.
- Rissanen J (1978) Modeling by shortest data description. *Automatica* 14:465–471.
- Rocchi F, Canella C, Noei S, Gutierrez-Barragan D, Coletta L, Galbusera A, Stuefer A, Vassanelli S, Pasqualetti M, Iurilli G, Panzeri S, Gozzi A (2022) Increased fMRI connectivity upon chemogenetic inhibition of the mouse prefrontal cortex. *Nat Commun* 13:1056.
- Rusina E, Bernard C, Williamson A (2021) Kainic acid models of temporal lobe epilepsy. *eNeuro* 8:ENEURO.0337-20.2021.
- Schneidman E, Berry MJ 2nd, Segev R, Bialek W (2006) Weak pairwise correlations imply strongly correlated network states in a neural population. *Nature* 440:1007–1012.
- Schreiber T (2000) Measuring information transfer. *Phys Rev Lett* 85:461–464.
- Shannon CE (1948) A mathematical theory of communication. *Bell Syst Tech J* 27:379–423.
- Suarez LM, Cid E, Gal B, Inostroza M, Brotons-Mas JR, Gomez-Dominguez D, de la Prida LM, Solis JM (2012) Systemic injection of kainic acid differentially affects LTP magnitude depending on its epileptogenic efficiency. *PLoS One* 7:e48128.
- Taatgen NA (2013) The nature and transfer of cognitive skills. *Psychol Rev* 120:439–471.
- Tait L, Tamagnini F, Stothart G, Barvas E, Monaldini C, Frusciantone R, Volpini M, Guttman S, Coulthard E, Brown JT, Kazanina N, Goodfellow M (2020) EEG microstate complexity for aiding early diagnosis of Alzheimer's disease. *Sci Rep* 10:17627.
- Tatum WO 4th (2012) Mesial temporal lobe epilepsy. *J Clin Neurophysiol* 29:356–365.
- Valero M, Averkin RG, Fernandez-Lamo I, Aguilar J, Lopez-Pigozzi D, Brotons-Mas JR, Cid E, Tamas G, Menendez de la Prida L (2017) Mechanisms for selective single-cell reactivation during offline sharp-wave ripples and their distortion by fast ripples. *Neuron* 94:1234–1247.e7.
- Van de Ville D, Britz J, Michel CM (2010) EEG microstate sequences in healthy humans at rest reveal scale-free dynamics. *Proc Natl Acad Sci USA* 107:18179–18184.
- Voges N, Hausmann J, Brovelli A, Battaglia D (2022) Decomposing neural circuit function into information processing primitives. *bioRxiv* 502783. <https://doi.org/10.1101/2022.08.04.502783>.
- Wenzel M, Han S, Smith EH, Hoel E, Greger B, House PA, Yuste R (2019) Reduced repertoire of cortical microstates and neuronal ensembles in medically induced loss of consciousness. *Cell Syst* 8:467–474.e4.
- Wibral M, Lizier JT, Vogler S, Priesemann V, Galuske R (2014) Local active information storage as a tool to understand distributed neural information processing. *Front Neuroinform* 8:1.
- Wibral M, Priesemann V, Kay JW, Lizier JT, Phillips WA (2017) Partial information decomposition as a unified approach to the specification of neural goal functions. *Brain Cogn* 112:25–38.
- Witter MP, Griffioen AW, Jorritsma-Byham B, Krijnen JL (1988) Entorhinal projections to the hippocampal CA1 region in the rat: an underestimated pathway. *Neurosci Lett* 85:193–198.
- Witter MP, Groenewegen HJ, Lopes da Silva FH, Lohman AH (1989) Functional organization of the extrinsic and intrinsic circuitry of the parahippocampal region. *Prog Neurobiol* 33:161–253.