Topological Analysis of Differential Effects of Ketamine and Propofol Anesthesia on Brain Dynamics

Thomas F. Varley[1,2], Vanessa Denny[1], Olaf Sporns[1,3], Alice Patania[3]

April 4, 2020

[1] Psychological & Brain Sciences, Indiana University, Bloomington, IN 47401
[2] School of Informatics, Computing & Engineering, Indiana University, Bloomington, IN 47401
[3] Indiana University Network Sciences Institute (IUNI), Bloomington, IN 47401

Abstract

Research into the neural correlates of consciousness has found that the vividness and complexity of conscious experience is related to the structure of brain dynamics, and that alterations to consciousness track changes in temporal evolution of brain states. Despite inducing externally similar states, propofol and ketamine produce different subjective states of consciousness. Here we explore the different effects of these two anaesthetics on the structure of dynamical attractors reconstructed from brain activity recorded electrophysiologically from cerebral cortex of a non-human primate. We used two different methods of attractor reconstruction. The first embeds the recordings in a continuous high-dimensional manifold on which we use topological data analysis to infer the presence (or absence) of higher-order dynamics. The second reconstruction, an ordinal partition network embedding, allows us to create a discrete state-transition network approximation of a continuous attractor, which is amenable to information-theoretic analysis and contains rich information about state-transition dynamics. We find significant differences between the awake, ketamine, and propofol conditions. Propofol significantly decreased the expression of higher-order structure in brain dynamics, while ketamine had the opposite effect, promoting an increase in complex dynamic features. Overall we found that ketamine produced dynamics similar to normal waking consciousness while propofol did not, and the significance of these findings is discussed.

Keywords: Anesthesia, consciousness, ketamine, propofol, attractor, dynamics

1 Introduction

In recent decades, the study of neural correlates of consciousness has developed into a rich and rapidly maturing field of research. A core component of the study of consciousness is the use of consciousness-altering drugs, which provide a mapping between measurable differences in brain dynamics and specific qualities of conscious experience [59]. Pharmacology has revealed the molecular actions of different drugs [29], and recent research has begun to look at comparing brain dynamics induced by

different drugs with a specific focus on how different dynamics might relate to conscious awareness [55, 63]. In this paper, we take such a comparative approach to explore the differences between the effects of propofol and ketamine on multi-scale brain dynamics with an eye specifically to how these dynamics might explain the differences in consciousness induced by both drugs. Here we characterize brain dynamics by adapting two complementary models capturing the evolution of whole-brain states through time: the first plots a trajectory through a high-dimensional configuration space, while the other discretizes transitions into a Markovian state-transition network.

While ketamine and propofol are both classified broadly as anesthetics, and both obliterate consciousness at high doses, it is useful to compare their different pharmacologies and the differences between the states they induce at low-to-moderate doses. By binding to $GABA_A$ receptors, propofol triggers widespread inhibition of neuronal activity and even at low doses, induces states of amnesia, sedation, atonia, and at higher doses, full anesthesia [29, 23]. In contrast, ketamine acts primarily as an antagonist of glutamaterigic NMDA receptors [29, 66], causing widespread, light central nervous system stimulation and a state typically referred to as "dissociative anesthesia" [16, 33]. Unlike propofol, which simply ablates consciousness, an otherwise unresponsive patient anesthetized with ketamine often continues to have complex, conscious experiences, including hallucinations, out-of-body experiences, and dream-like, immersive experiences [66].

It has been hypothesized that the differences between ketamine and propofol anesthesia are the result of changes to the global brain dynamics induced by each drug [63]. Experimental evidence suggests propofol inhibits the ability of the brain to maintain high levels of dynamical complexity, resulting in a low-entropy state insufficient for supporting conscious awareness [54, 52]. In contrast, ketamine's blockade of NMDArs is thought to dis-inhibit cortical neurons, causing widespread, uncoordinated excitatory activity [66, 55, 63]. This may result in an increase in the entropy of brain activity without abolishing consciousness, artificially expanding (or at least altering) the state-space repertoire. The hypothesis that a dynamic state of higher-than-normal entropy might correspond to a psychedelic or hallucinatory state of consciousness has become known as the Entropic Brain Hypothesis [6, 5] and received empirical support from studies of psychedelic drugs [4]. The majority of these studies rely ultimately on point-summaries of "complexity" (eg: Lempel-Ziv complexity [54, 52], fractal dimension [34, 61], entropy [36, 68], etc). However, these point-summary measures, while informative, collapse multi-scale dynamics into a single number and thus have difficulty capturing its specific shape or form.

We report results relating brain dynamics to states of consciousness (normal wakefulness, ketamine anaesthesia and propofol anaesthesia) using freely-available electrophysiological data from the NeuroTycho project [42, 65]. In the first part, we describe the structure of brain activity in terms of its evolution through a high-dimensional state-space: at every moment t_i, the system's "state" can be described as a vector embedded in a k-dimensional space, one dimension for each channel. Conveniently, this embedded point cloud (EPC) does not require any dimensionality reduction and conceivably, all other metrics can be reconstructed from this high-dimensional structure. One of the simplest measure is the "distance" between successive points, from which we can derive a "velocity", the distance traversed between subsequent time-steps, as well as reconstruct how "far" in state-space the system moves over the course of the recording and how these high-dimensional dynamics vary

between states. In addition to these local measures, modeling the EPC as a manifold renders it amenable to techniques from algebraic topology (a branch of mathematics dealing with the structures of surfaces and manifolds, particularly focused on homology). Topological data analysis [9, 56], or TDA, allows us to understand details of the trajectory, including the emergence of cycles and what regions it may preferentially visit. TDA has been used extensively to characterize the chaoticity of time series [11, 48, 64, 46] and provides a suite of techniques for classifying high-dimensional structures.

One limitation of the EPC approach, is that the data is represented by continuous variables, and so every "state" is unique, making it difficult to understand higher-level state-transition dynamics. To address this issue, in the second part of the study we discretize the EPC by creating ordinal partition networks [57, 38, 41] (OPNs) which map sets of multiple unique states to the same set of nodes in a network and for which the probability of the system transitioning from Set A to Set B is recorded as a weighted directed edge from node A to node B. In this way, the OPN represents a reconstructed attractor that is discrete rather than continuous and consequently amenable to a number of analyses specific to discrete manifolds like networks. In particular, network science provides tools that allow us to explore the attractor at different levels: the micro-scale (node-level differences), meso-scale (community-level differences), and macro-scales (global topological differences). In combination, the two parts of our study allow a much richer understanding of brain dynamics than single-point measures can provide.

2 Materials & Methods

2.1 Data

We used the Neurotycho dataset, an open-access collection of multidimensional, invasive electrocorticographical recordings from multiple macaque monkeys [42]. Specifically, we used data from two monkeys (Chibi and George), prior to, and during anesthesia with propofol or ketamine [65]. ECoG data was recorded on 128-channel, invasive recording array at a sampling rate of 1000 Hz. For both monkeys, the array covered the entire left cortical hemisphere, including the medial wall.

2.1.1 Anesthetic Induction

The details of anesthetic induction are described in [65], and can be viewed on the relevant Neurotycho wiki page[1]. Briefly, in the Awake condition, the monkeys were restrained in a primate chair with arms and legs fixed, and neural data was collected while the monkey was calm. During the Ketamine anesthesia condition, the restrained monkeys were injected with intra-muscular ketamine (4.3mg/kg for Chibi, 5.9mg/kg for George), and anesthesia was determined as the point that the monkey no longer responded to physical stimulus (manipulation of the hand and/or tickling the nose with a cotton swab) and slow-wave oscillations were observed in recorded data. The data were then recorded for ten minutes (no supplemental or maintenance doses were given). For the Propofol condition, both restrained monkeys were injected with intra-venous propofol (5.2mg/kg for Chibi,

[1] http://wiki.neurotycho.org/Anesthesia_and_Sleep_Task_Details

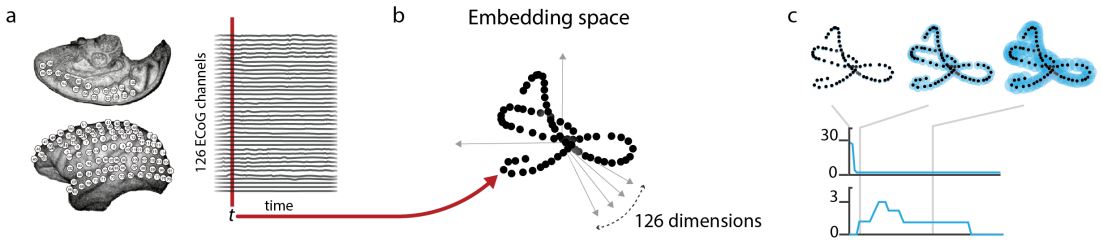

Figure 1: Construction of the embedded point cloud from a multidimensional time series. Every moment t in the time series corresponds to a multi-dimensional vector in a high-dimensional space (a), creating a point-cloud that traces out the trajectory of the system through this space as it evolves (b). We then computed the persistent homology of the complex defined by these points using cosine distance recording the evolution of connected components and cycles throughout the filtration (c). The process can be thought of as spheres centered on each embedded point with growing radius. As the radius of the spheres grow and begin to overlap, a discretized approximation of the underlying manifold emerges.

5mg/kg for George), and loss-of-consciousness assessed using the same criteria. Recordings were then carried out for ten minutes.

For each drug condition, recordings of normal consciousness were made prior to infusion of anesthesia, and each experiment was done twice (eg: two experiments where Chibi is anesthetized with propofol, two with ketamine, resulting in four recordings of Chibi awake, and two of Chibi under each condition). We removed channels with intractable artifacts from the scans (four channels from Chibi, two from George). This meant that the number of channels in the final analyzed data differed between monkey (124 channels for Chibi and 126 channels for George). From each recording we manually selected six, two-second intervals from the anesthetized condition and three two-second intervals from the Awake condition (because each anaesthesia condition had its own Awake condition, we had twice as much Awake data, and consequently halved the number of sections taken from each scan to ensure that the total number of Awake intervals matched the total number of anesthetized intervals for each drug condition). All selections were free of major artifacts. This resulted in 24 samples each for the Awake condition, the Ketamine condition, and the Propofol condition. The time series were visualized and manually selected in EEGLab [15].

2.1.2 Pre-Processing

After subsections of the time series had been selected, pre-processing was performed in MNE-Python [19, 20]. Each series was high-pass filtered with a low frequency of 0.5Hz, low-pass filtered with a high frequency of 200 Hz and notch-filtered at 50Hz and all subsequent harmonics up to 250Hz to account for electrical line noise in Japan, we removed 50ms from the start and the end of each sample due to filtering artifacts. All filters were FIR type and were applied forwards and backwards to avoid phase-shifting the data [22]. Following the filtering, time series from each sample were z-scored as a set (to relative distances between each moment).

2.2 Embedded Point Cloud

The analysis explored the distribution of fine-scale global activity patterns that appear over the course of the recording. In our framework, global activation patterns recorded by the array are represented as moving states in a high-dimensional configuration space. Studying the relative position and dynamics of these states gives unique insights in quantifying the underlying dynamical structure.

Each time segment can be represented as a $C \times T$ matrix M, where C is the number of channels (126 or 124), and T is the number of samples taken over the course of the recording. Each of the T column-vectors in M (each C entries long) represents a distribution of voltages across the recording array, at each moment in time.

We construct the embedded point cloud (EPC) by treating each column as a vector in C-dimensional space, so at each time stamp the brain state is encoded as a C-dimensional vector in the space of all possible states it can assume. As time goes by, we can see the entire recording as a trajectory across this continuous state space. The first analysis we did was to calculate the cosine distance between temporally consecutive states. Since the data was recorded with a constant sampling interval, the distance between states at time t and $t+1$ are proportional to a "velocity" through the state-space. The average velocity is a measure of how rapidly patterns of activity observed over the array change and the rate at which the system as a whole is evolving.

2.3 Topological Data Analysis

Topology is the area of mathematics that studies shapes and spaces hard to represent visually. More recently, topology has been applied in data analysis to help describe and classify noisy or high-dimensional data, for example by extracting topological invariants present in complex data sets [7, 44, 46]. Among the many topological methods that have been developed for data analysis, the most frequently used is persistent homology. Persistent homology allows us to build descriptors of the shape of a point cloud, by cataloguing the existence of different structural features, such as connected components, cycles, voids, etc. at different levels of coarse discretization of the data. We can think of the process (known variably as a Rips filtration or a Vietoris-Rips filtration) as spheres centered on each embedded point with growing radius. As the radius of the spheres grows, they begin to intersect, connecting the data points, and a representation of the shape of the data will start to emerge. As one increases the radius, at what scale do we observe changes in representation of the data? To quantify this changing representation, persistent homology takes note of the number of connected components (0-dim homology) and the radius at which they merge together and when the points in these components, start to connect, creating more complex structures like loops (1-dim homology), and at what size of the spheres they disappear.

To compare the features, we use Betti curves. For each feature type (connected component and cycles) a Betti curve counts the number of features that exist at each scale in the Rips filtration. Since the curves all have the same support, the increasing radius of the spheres in the embedding space, the curves can be compared between each other or averaged together, finding the radius that create the most complex feature sets.

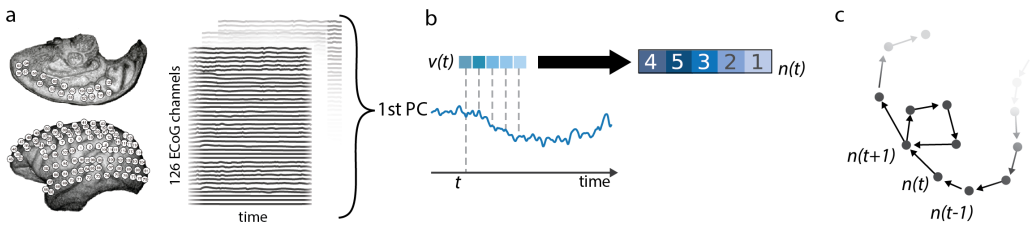

Figure 2: We constructed the OPN of the projection of each 2s interval multi-variate time series onto the first temporal principal component of the entire data set (a). An OPN was computed for each projection separately. Each PC time series is embedded in a d dimensional space, using a time-lag τ, as is done when constructing a Taken's Embedding. The result is a temporally ordered set of vectors of length d, v_t. Each vector v_t is then mapped to the permutation, π which sorts the coefficients of the vector in increasing order (b). This new vector $n_t = \pi(v_t)$ will be represented by a permutation of the numbers $1, 2, \ldots, d$. We can consider these permutations as nodes in a directed network and connect with a directed edge two permutations that come from consecutive time points (c).

2.4 Ordinal Partition Network

Looking at permutation sequences in a time series has been a rising trend in the dynamical systems community in the last 20 years [1, 3]. Permutations are a sensitive indicator of the dynamic state of a system and can be efficiently computed, even for long time series data. One significant benefit of constructing permutations is that it maps a continuous time series to a finite set of discrete permutations, which allows for principled information-theoretic analysis of systems that might not otherwise be amenable, such as permutation entropy [1].

To explore the temporal dynamics of the system, we constructed ordinal partition network (OPN) representations of the data [57, 38], following the procedure discussed in [41]. Due to the difficulties associated with multi-dimensional OPNs [67], we applied this method to a uni-variate time series, constructed from the first temporal principle component of the data. For each monkey, we concatenated all 36 time segments and using PCA, projected them onto a shared manifold. By projecting all samples into the same space, we ensure that shared dynamics are comparable across scans. Projecting every multi-variate time segment onto the first principal component, we obtained a single uni-variate time series containing most of the dynamical information about the global recorded signal, which we then cut back into the 36 relevant slices (see Figure 2). By using only the first temporal principal component, we can capture the maximum amount of dynamic variance in a single time series.

To construct an OPN, begin with a time series $X = x_1, x_2, ...x_n$. This time series is then embedded in a d dimensional space, using a time-lag τ, as is done when constructing a Taken's Embedding. The result is a temporally ordered set of vectors of length d, where each $v_i = [x_i, x_{i+\tau}, \ldots, x_{i+(d-1)\tau}]$. Each vector v_i is then mapped to the permutation, π which sorts the coefficients of the vector in increasing order. Sorting the coefficients of v_i we will have that $v_{i,1} \leq v_{i,2} \leq \cdots \leq v_{i,d}]$. Each coefficient is then replaced by the position they have in this ordering π. This new vector will be represented by a permutation of the numbers $1, 2, \ldots, d$ $n_i = [\pi(x_i), \pi(x_{i+\tau}), \ldots, \pi(x_{i+(d-1)\tau})]$. We can consider these permutations as nodes in a directed network and connect with a directed edge

two permutations that come from consecutive time points. The resulting transition network will have less nodes than existing time points in the original time series, as there might be multiple i for which the respective delay vectors v_i give the same permutation π. To incorporate this information into the resulting OPN, we give a weight to each node in the network counting how many time points i led to the same permutation.

In addition to the topological and information-theoretic measures reported here, we ran a battery of more standard network measures aiming to characterize the connectivity of the OPNs, including measures of centrality (betweenness, Katz), clustering coefficient, and modularity (using the Infomap algorithm [50]. For results and discussion see the Supplementary Material (Fig. S.I. 1).

2.4.1 Free Parameters

Like other methods of attractor reconstruction, the OPN algorithm requires two free parameters: the embedding dimension d and the time-lag, τ. There is no agreed-upon method for choosing the optimal d and τ. Different researchers have suggested different criteria, including values that maximize the variance of the degree distribution [38], false-nearest-neighbours criteria [28], or the first local maxima in the permutation entropy [41]. One significant consideration is that the embedding dimension must be large enough that the state-space is sufficiently large to fully capture the range of patterns present in the data (otherwise the resulting OPN fails to capture all the dynamics), but not so large that every moment is unique (in which case, we reconstruct the time series as a path graph). The constructed network should have a sufficiently complex topology so as to represent the richness of the initial data set, while still finding meaningful patterns. See Figure S.I. 5 for an example on the effect of these parameters on the structure of an OPN.

Usually for finding the optimal lag in 1-dim time series, one of the parameters to look at is auto-correlation. We estimated the auto-correlation in the EPC as the size of the cliques formed when constructing the Vietoris-Rips complex at the radius at which all embedded points are connected. We use this estimate of the optimal lag for the OPNs. For choosing the optimal dimension we used both nearest neighbour statistics [31] of the embedding on the first principal component of the dataset and the percentage of states not in common between awake and anesthetized condition, see Figure S.I. 4. Both these methods agreed on an optimal dimension of $d = 7$.

2.4.2 Information-Theoretic Analysis of OPNs

A significant benefit of the OPN is that, being a discrete manifold, the state-transition network is amenable to information-theoretic analysis in a way that a continuous manifold is not. Here we report three measures of how much information a system encodes in its state transition graph [25, 24, 30]. The first measure is the determinism, which measures of how much information a system's state transition graph encodes about its future evolution; i.e. how deterministic, on average, is the evolution from state i to state j. The determinism is low in a system where each state has an equiprobable chance of evolving into one of many future states, while a system where each state evolves with probability equal to 1 to a subsequent state would be highly deterministic. The average determinism of a directed network X with N can be quantified as:

$$Det(X) = \frac{log_2(N) - \langle H(W_i^{out})\rangle}{log_2(N)}$$

Where $\langle H(W_i^{out})\rangle$ corresponds to the average entropy of the probability distribution of possible futures (weighted out-going edges) for each node i. The second measure is the degeneracy, which gives a measure of how much information a system's state transition graph encodes about its past evolution. A system where all states feed into the same future would be described as highly degenerate, while a system where each state had a well-defined past would exhibit low degeneracy. Degeneracy is calculated as:

$$Deg(X) = \frac{log_2(N) - H(\langle W_i^{out}\rangle)}{log_2(N)}$$

From the determinism and the degeneracy we can calculate the effective information of the graph, which is a measure of the total amount of information the system encodes in its causal states (for our analysis, we normalize determinism, degeneracy, and effective information by $log_2(N)$ so as to control for variability in network sizes). Effective information is calculated as:

$$Eff(X) = Det(X) - Deg(X)$$

These three measures constitute an information-theoretic analysis of state-transition graphs representing a system outputting a continuous signal. The OPN can be thought of as a realization of the idea of an ϵ-machine [12, 13], which provides an optimal approximation of a dynamical system based on output data. Combining this information-theoretic formalism provided by Hoel et al [25, 30]. with the OPN formalism provides a computationally tractable set of tools to explore the informational structure encoded in continuously-varying signals.

2.5 Software

All statistical analysis was carried out in Python using the Scipy Stats package (version 1.1.0) [43, 39]. Analysis of variance was computed using the Kruskal-Wallis Analysis of Variance test [32], and post hoc testing was done using the Mann-Whitney U test [37]. Non-parametric statistics were chosen due to the uncertainty that the data were sufficiently Gaussian. Persistence homology analysis was done using the Ripser (version 0.3.2) [58] and Persim libraries (version 0.0.9) as part of the Scikit-TDA library (version 0.0.4) [53]. Other packages used include the Numpy library (version 1.15.4) [62], Scikit-Learn (version 0.20.0) [45], Matplotlib (version 2.2.2), [26], Spyder (version 3.2.3), NetworkX (version 2.2) [21], iGraph (version 0.7.1) [14]. Analysis was done in the Anaconda Python Environment (Python 3.6.8, Anaconda version 5.0.0).

3 Results

3.1 Topological data Analysis of the Embedded Point Cloud

To explore the distribution of patterns of activity across the cortex, we embedded each moment in time (which corresponds to a high-dimensional vector according to the number of channels in the original data) in space. As a preliminary analysis on the global dynamical patterns represented by the embedded points, we compared the distance in the embedded space between consecutive points to quantify the change of the global brain states in 1 ms steps (the sampling rate of the data). We found that the rate of change is higher in the Awake condition than in both the anesthetized ones, suggesting a lower rate of variability of brain states in the anesthetized monkey than the Awake (see Fig.3 a). This suggests that, when awake, the brain is traversing a "larger" space of possible states and traveling through them faster than it does when anesthetized (and is restricted to exploring a smaller subspace of states). This measure is very local and relative to each point in time individually. To have a better understanding of the dynamic as a whole we need to look at all the time points at once and their relationship in this space. To analyze the entire spectrum of global brain states dynamics in the different conditions, for each 2s time interval studied we ran a Vietoris–Rips filtration on the resulting embedded point cloud. The results describe the distribution of higher-order patterns that appear in each state.

Because the embedded point clouds all exist in the same configuration space, we compare the results not only by the number of homological features detected, but also the radius necessary for these features to appear in the filtration. We found that states in the Awake condition are further apart, and one has to connect states at larger distances to obtain a single connected component (Fig 3 c). This minimal distance necessary to have one emergent component decreases significantly in both anaesthetised states. These results mirror what we found when examining the velocity of the system trajectory, suggesting that during the Awake condition the dynamics reaches states that are very different from each other (consequently, farther apart in configuration space). While the velocity only captures consecutive activation patterns, here we see that this pattern is consistent across all nearest neighbors needing smaller radius spheres to fully connect a single time-segment embedded point cloud with ketamine and further more with propofol infusions.

The emergence of cycles in Vietoris-Rips filtration has been associated to both a signature of the local geometry (small cycles) [2] and of global topology (or shape) [7] of a data set. The local geometry in this case can be interpreted as local noise or disturbance in the way the brain transitions between states (short-lived cycles), while global topology (or shape) of an embedded time series has long been associated with the detection of chaotic, periodic and quasi-periodic behaviour of the underlying dynamic [40, 27, 48]. From the Betti curves shown in Fig 3e we notice a general decrease in the number of cycles in anesthetized compared to Awake conditions implying a loss of complexity in the global brain dynamic patterns detectable in the EPC (subject level analysis of the Betti Curves can be found in Figure S.I. 2). An interesting feature is the increase in short lived cycles at low filtration values during the ketamine sedation, which hints at the increased presence of high frequency noise in the temporal path as hypothesized in the literature. To validate these results we studied some standard homological measures. Analysis of variance of the lifetime of the longest-

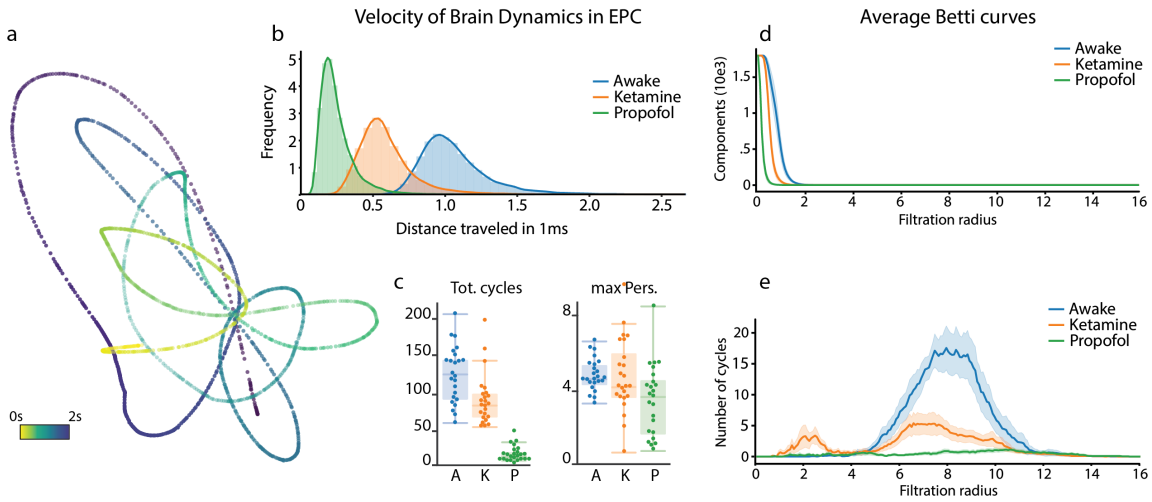

Figure 3: Results for the EPC analysis of brain dynamical trajectories in the high-dimension configuration space. (a) A 3-dimensional embedding of the 126-dimensional time series of a representative time segment of the Awake condition. (b) Empirical distributions of the instantaneous velocity for the three conditions. (c) The total number of homological cycles and life duration of the longest cycle in each time segment (represented by a point in the swarm plots), separated by condition (A = Awake, K = Ketamine, P = Propofol). We can see here that the Awake condition had on average a significantly higher total number of cycles (124.3 ± 37.05) than either the Propofol condition, which had 14.58 ± 9.96 cycles, or the Ketamine condition which had 89.79 ± 33.9 cycles. Moreover, the Awake condition had on average the longest maximum persistence (4.76 ± 0.83), followed by the Ketamine condition (4.72 ± 1.88), and the Propofol condition had the smallest (3.44 ± 1.82). (d-e) Avarage Betti curves, with a 90% confidence interval, counting the number of homological features (connected components (c) and cycles (d)) present in the complex as the spheres radius increases.

persisting cycles found slight differences between the states ($H = 8.93, p = 0.01$). Post-hoc testing found significant differences between the Awake and Propofol conditions ($U = 61, p = 0.01$) and the Ketamine and Propofol conditions ($U = 79, p = 0.04$), but not between the Awake and Ketamine conditions. The Awake condition had on average the longest maximum persistence (4.76 ± 0.83), followed by the Ketamine condition (4.72 ± 1.88), and the Propofol condition had the smallest (3.44 ± 1.82). This strongly suggests that Propofol is disrupting the formation of higher-order dynamics in a way that the Awake condition and Ketamine condition are not.

We also collected the total number of cycles to appear over the course of the entire filtration and the maximum number of cycles that exists at any point in the filtration. Both of these had a similar pattern across states. Analysis of variance found significant differences between all conditions in the total number of cycles ($H = 51.87, p = 4.43 \times 10^{-12}$). The Awake condition had on average a significantly higher total number of cycles (124.3 ± 37.05) than either the Propofol condition ($U = 0, p = 1.81 \times 10^{-5}$), which had 14.58 ± 9.96 cycles, or the Ketamine condition ($U = 49.5, p = 0.007$) which had 89.79 ± 33.9 cycles. Ketamine had a significantly higher value than the Propofol ($U = 0, p = 1.81 \times 10^{-5}$), which was the smallest.

There were also significant differences between the states in the maximum number of cycles to

appear at a single moment ($H = 58.56, p = 1.93 \times 10^{-13}$). Here, the trend was similar to the maximum persistence: the Awake condition (27.04 ± 10.86) was significantly higher than either the Ketamine ($U = 1.5, p = 4.91 \times 10^{-5}$) or the Propofol conditions ($U = 0.0, p = 1.8 \times 10^{-5}$). The Ketamine condition (12.17 ± 4.47) was significantly higher than the Propofol condition ($U = 0.0, p = 1.81 \times 10^{-5}$), which was the lowest (3.08 ± 1.35). As before, the most significant pattern is the collapse of cycles in the Propofol condition (in this case, almost an order of magnitude fewer cycles than in the Awake condition). This indicates that the change in cyclic dynamics in the Propofol condition is a multi-scale phenomenon, an intuition born out by the analysis of the Betti curves in Figure 3.

3.2 Ordinal Partition Network

3.2.1 Basic Network Structure

Simple measures of the OPN provide relevant information about the structure of the dynamic attractor and the behavior of the brain in different states. Arguably, the simplest is the number of nodes in the OPN, which corresponds to the number of uniquely realized states in the time series. Kruskal-Wallis analysis of variance found significant differences between all three conditions (H=31.37, p=1.78×10^{-7}). Post-hoc analysis found that there was a significant reduction in the number of nodes between the Propofol condition (456.25 ± 57.1) and both the Awake ($U = 19, p = 0.00018$) and Ketamine ($U = 17, p = 0.00014$) conditions, but not between Awake (563.0 ± 68.0) and Ketamine (615.875 ± 110.5). This finding is consistent with previous findings that Ketamine, unlike Propofol, maintains or increases the algorithmic complexity of brain activity [55, 63].

Another simple but informative measure is the variance in how frequently all nodes are visited (this is related to, but separate from, the notion of degeneracy described in section 2.4). If all nodes are visited roughly equal numbers of times, the variance will be low and the system can be thought of as reasonably "flat", lacking variability in the recorded dynamical emerging patterns. In contrast, if some nodes are visited very frequently compared to others, they may represent attractor states that the brain repeatedly returns to. Analysis of variance found significant differences between all three conditions on the variance of the number of times nodes were visited (H=35.5, p=1.9×10^{-8}). As with the number of states, post-hoc analysis found that there were no significant differences between the Awake (135.7 ± 80.27) and Ketamine (97.0 ± 96.5) conditions, but significant differences between both of those conditions and Propofol (340.1 ± 127.6). These results suggest that the Propofol condition has a subset of states that it frequently returns to after rare excursions, while the awake and the ketamine conditions visit states more equitably.

To further characterize the state-transition dynamics, we also calculated the proportion of states that were only visited once. Related to the variance in the number of visits, which describes attractor dynamics, the percentage of nodes only visited once gives information about the comparative flexibility of the system. We found significant differences between all three conditions (H=15.62, p=0.0004). Post-hoc analysis found that the condition with the greatest percentage of nodes in the network with a single visit was the Awake condition (61.9% ± 3.0%), which was significantly higher than the Ketamine condition ($U = 39, p = 0.0015$), but surprisingly, not the Propofol condition. The

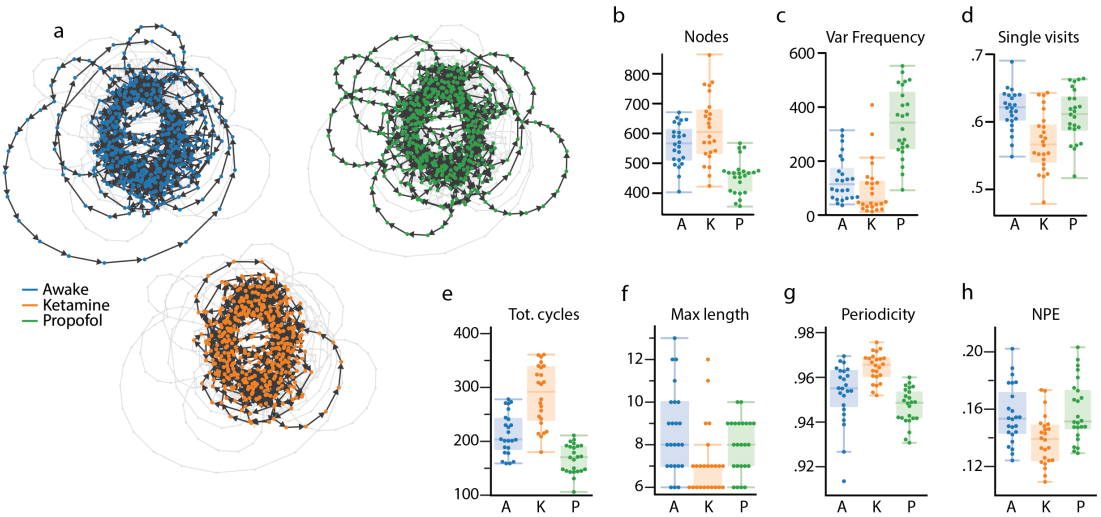

Figure 4: Results of the OPN analysis. (a) Representation of three representative OPN networks, one per condition. (b-d) Summary statistics for structural measures of the OPN networks (each represented by a data point in the swarm plots), separated by condition (A = Awake, K = Ketamine, P = Propofol). (b) There was on average a significant reduction in the number of nodes between the Propofol condition (456.25±57.1) and both the Awake (563.0±68.0) and Ketamine (615.875±110.5) condition. (c) Similarly on average there were no significant differences in the variance of the number of times nodes were visited between the Awake (135.7±80.27) and Ketamine (97.0±96.5) conditions, but significant differences between both of those conditions and Propofol (340.1 ± 127.6). (d) The condition with the greatest percentage of nodes in the network with a single visit was the Awake condition (61.9 ± 3.0). The Propofol condition had the next higher proportion of single visits (60.89 ± 3.8), followed by the Ketamine condition (57.0 ± 4.3). (e-h) Summary Statistics of the OPN cyclic structure in the OPN (each represented by a data point in the swarm plots), separated by condition (A = Awake, K = Ketamine, P = Propofol). (e) The Ketamine condition had the highest number of cycles (285.125 ± 54.93), which was significantly higher than the Awake condition which had 213.58 ± 37.85 cycles and the Propofol condition, which had 168 ± 26.5 cycles. (f) (g) Ketamine had the highest periodicity score (0.96 ± 0.006), followed by the Awake condition (0.95 ± 0.013), with Propofol having the smallest (0.947 ±0.008). (h) Propofol had the highest normalized persistence entropy (0.16 ± 0.02), followed by the Awake condition (0.16 ± 0.02), and Ketamine had the smallest (0.14 ± 0.017).

Propofol condition had the next higher proportion of single visits (60.89% ± 3.8%), followed by the Ketamine condition (57.0% ± 4.3%). A productive future reseach project would be to characterize how these various measures relate to each other: which measures are redundant, and which are synergistic? We might have hypothesized that conditions like Propofol, which visit the same states over and over again would also have a lower percentage of states visited once, however this does not appear to be the case.

3.2.2 Topological Analysis

Previous work by Myers et al., [41] showed that persistence homology of OPNs can be informative about the dynamics driving a complex system. As with the EPC, the presence of cycles and higher-order structures in the discrete state-transition dynamics can shed light on emergent properties specific to the conditions. To explore this domain, we performed TDA on the OPNs which allowed us to list all closed directed paths (loops). This analysis allows us to study how often the system revisits past states, and how many states are visited before this occurs.

The first measure we studied was the total number of loops in each OPN. Analysis of variance found significant differences between the three conditions (H=42.1, p=7.22×10^{-10}). Post-hoc testing found that the Ketamine condition had the highest number of cycles (285.125 ± 54.93), which was significantly higher than the Awake condition ($U = 20, p = 0.0002$), which had 213.58 ± 37.85 cycles and the Propofol condition ($U = 1, p = 2 \times 10^{-5}$), which had 168 ± 26.5 cycles.

We also used two slightly more complex TDA measures, the normalized persistence entropy [41] and the periodicity [47]. Analysis of variance found a significant difference between conditions (H=12.4, p=0.002), and post-hoc analysis found significant differences between the Propofol condition and the Ketamine condition ($U = 63, p = 0.013$) as well as between the Awake and Ketamine conditions ($U = 65, p = 0.015$). Propofol had the highest normalized persistence entropy (0.16 ± 0.02), followed by the Awake condition (0.16 ± 0.02), and Ketamine had the smallest (0.14 ± 0.017). This suggests that the distribution of cycles is restructured under Ketamine in a way that it is not in the Awake or Propofol conditions.

The final measure we tested was the periodicity score [47] which relates size of the longest loop relative to the size of the OPN. Networks with a loop that includes all of its nodes will be highly periodic, as there is a path that traverses every state of the system in a periodic manner. The shorter the longest loop, the less can be said about a single periodic dynamic driving the transitions. Analysis of variance found significant difference between the conditions (H=32.9, p=1.17^{-8}). Post-hoc testing found a pattern similar to the normalized persistence entropy, with significant differences between the Awake and Ketamine conditions ($U = 35, p = 0.001$) and between the Ketamine and Propofol conditions ($U = 0, p = 1.8 \times 10^{-5}$, but no significant differences between Awake and Propofol. Ketamine had the highest periodicity score (0.96 ± 0.006), followed by the Awake condition (0.95 ± 0.013), with Propofol having the smallest (0.947 ±0.008). Both TDA measures on the OPN found patterns that indicate similarity between the Awake and Propofol conditions, suggesting that this representation captures dynamics that are distinctly induced by Ketamine but preserved under Propofol anesthesia.

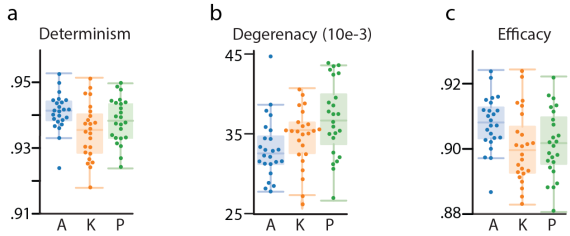

Figure 5: Summary statistics of three different information-theoretic characterizations of the OPN change by condition. (a) The Awake condition had the highest determinism (0.94 ± 0.006), significantly higher than the Ketamine condition, which had a mean determinism of 0.935 ± 0.008, but not significantly different from the Propofol condition (0.938 ± 0.007). (b) The Propofol condition had the highest degeneracy (0.037 ± 0.004), which was significantly higher than the Awake condition ($U = 106, p = 0.009$), but not the Ketamine condition. Ketamine had the next highest degeneracy (0.034 ± 0.0036), which was not significantly different from the Awake or Ketamine conditions. The Awake condition had the lowest degeneracy (0.033 ± 0.0037). (c) Since effectiveness is defined as the difference between the determinism and the degeneracy, and the determinism in all cases was near maximal (> 0.9). The Awake condition had the highest effectiveness (0.91 ± 0.008). The Propofol and Ketamine conditions had functionally the same mean effectiveness (0.9 ± 0.01) and consequently were not significantly different.

3.2.3 Information-Theoretic Analysis

Kruskal-Wallis analysis of variance found significant differences between the determinism of all three conditions (H=7.1, p=0.029), the degeneracy of all three conditions (H=7.95, p=0.018), and the effectiveness (H=7.68, p=0.022). For visualization see Figure 5. The Awake condition had the highest determinism (0.94 ± 0.006), significantly higher than the Ketamine condition ($U = 76, p = 0.034$), which had a mean determinism of 0.935 ± 0.008, but not significantly different form the Propofol condition (0.938 ± 0.007). There was no significant difference between the Ketamine and Propofol conditions.

Unsurprisingly, the Propofol condition had the highest degeneracy (0.037 ± 0.004), which was significantly higher than the Awake condition ($U = 106, p = 0.009$), but not the Ketamine condition. Ketamine had the next highest degeneracy (0.034 ± 0.0036), which was not significantly different from the Awake or Ketamine conditions. The Awake condition had the lowest degeneracy (0.033 ± 0.0037). Since effectiveness is defined as the difference between the determinism and the degeneracy, and the determinism in all cases was near maximal (> 0.9), the results of the effectiveness calculation are largely dominated by the determinism and followed the same pattern. The Awake condition had the highest effectiveness (0.91 ± 0.008), which was significantly higher than both the Ketamine condition ($U = 81, p = 0.049$) and the Propofol condition ($U = 0.017, p = 0.017$). The Propofol and Ketamine conditions had functionally the same mean effectiveness (0.9 ± 0.01) and consequently were not significantly different.

These results show how different drugs can exert distinct effects on the information dynamics of the brain. While both Ketamine and Propofol had significantly lower effectiveness compared to the Awake condition, when decomposed into determinism and degeneracy, we can see that they are driven by different changes. The Ketamine condition had the lowest determinism (the most uncertainty

about the future of the system), while Propofol had the highest degeneracy (the most canalizing dynamics, where multiple distinct states feed into the same future state). Both patterns result in loss of effective information but for distinctly different reasons and indicate different underlying state-transition dynamics.

4 Discussion

The two methodological approaches described here (topological data analysis, and ordinal partition networks) belong to an emerging family of techniques for analyzing univariate and multivariate time series by representing them as a complex network [69]. Each approach constructs different networks that encode different aspects of system dynamics into their topology. The embedded point cloud/TDA maps distinct spatial patterns of activity across the cortex to nodes, and the similarity between them to undirected edges, while the OPN maps distinct temporal patterns to nodes and the probability of subsequent patterns to weighted, directed edges. By bringing these methods together, we can create a more holistic portrait of brain dynamics and more fully explore how they relate to consciousness than a single univariate measures. While single-point measures have shown utility in identifying changes to conscious awareness [8, 52, 54, 55, 60, 61], a multidimensional analysis can supplement these results by providing a more detailed breakdown of which features of dynamics are significant.

In this study, we found that, while ketamine and propofol induce states that are largely indistinguishable to external observation (recall that the same behavioral criteria was used for each drug to determine loss of consciousness), they produce distinct dynamical regimes, which differ from each other and waking consciousness along several axes. The benefit of this kind of multi-dimensional analysis are obvious: by finding the ways in which the dynamics are different, and the ways in which they are similar, we can begin to tease apart which dynamical characteristics may be necessary to support particular aspects of behavior and cognition.

While each state is unique, the relative similarities between states seemed at least partially dependent on the analysis being run. In general, it seems that the Awake condition has more in common with the Ketamine condition than the Propofol condition. For instance, TDA on the EPC found smaller differences between the Awake and Ketamine conditions for two measures (for instance, the total number of cycles and the maximum number of cycles), and no significant difference between Awake and Ketamine for the maximum persistence. Similarly, the two most basic OPN measures (number of nodes, variance in frequency of visits) indicated that the Awake and Ketamine conditions were more similar, while the Propofol condition was something of an outlier.

Interestingly, this pattern was not preserved when doing TDA on the OPNs: for three of the four measures (number of cycles, maximum cycle length, and periodicity score), it was the awake and the propofol conditions that were more similar, and the ketamine condition that was the odd one out (although significant differences were found between all three most times). The exact significance of this is unclear, however it suggests that the TDA of OPNs captures dynamical relationships between the different states that TDA of the EPC, or the standard OPN analyses do not. The work by Myers et al. [41] showed that TDA of OPNs can retrieve significant information about the

system dynamics such as Lyapunov exponent even in the presence of noise and so future research aimed at understanding the applicability of these analyses to neural data look to be promising. If the similarities between the dynamics produced by ketamine and normal awareness are indicative of phenomenological consciousness, the similarities between normal awareness and propofol may be indicative of ketamine's specific effect on NMDA receptors that are not shared by propofol [66].

The question of how to interpret cycles in the context of neural data is an open one, both in the EPC and the OPN. One possible interpretation is that cycles can be thought of as "higher order dependencies" involving multiple channels. Rosas et al., describe patterns as "deviations from independence" [49], suggesting that an EPC with no cycles is one where all dimensions (i.e. channels) operate independently of one another through time. In the context of neural data of this type, if all channels were independent, the EPC would tend towards a 126-dimensional Gaussian distribution with maximal entropy, approximating a high-dimensional sphere, presumably without any significant or long-lived cycles. Patterns, such as cycles and voids, may then be interpreted as deviations from the maximally entropic configuration.

This may be one explanation for the collapse of cyclic structures in the EPC from the Propofol condition: if all channels behave independently (over the relevant time-scales), higher-order strutures cannot emerge. This would be consistent with the hypothesis that propofol exerts its consciousness-altering effects by disrupting coordination between brain regions [35] - in this context, disrupted communication may appear as a collapse of higher-order structural dynamics that would ordinarily recruit these channels. One possible future direction would be to relate changes in the structure of embedded point clouds to more standard measures of multi-region coordination, such as functional connectivity. In the temporal domain, propofol has a clear effect of reducing the number of unique temporal states visited, which is consistent with previously reported results [68, 54, 63, 61]. In addition, there is an apparent emergence of powerful attractor states that the brain repeatedly visits, over and over again, in contrast to both normal waking consciousness and the ketamine state. These results are consistent with the entropic brain hypothesis and suggest that loss of consciousness (as opposed to dissociation of consciousness) involves a "deepening" or "canalizing" of the landscape of the brain's state-space [5]. This is quite relevant to perturbation-based measures of consciousness, such as the Perturbational Complexity Index (PCI), [8, 52]: if it is true that loss of consciousness can be thought of as canalizing the state-space, we would expect that perturbations of the brain's state should persist for shorter periods of time during loss of consciousness than when subjects are conscious. Future work revisiting the effect of TMS in propofol and ketamine may shed light on this possibility. This could also be a useful diagnostic criteria for disorders of consciousness: recent work has generalized the original PCI measure [10, 51] and relative canalization may be a useful measure to include.

The Rips filtration can be understood as being closely related to the structure of recurrence plots, commonly used in chaos theory [18], although here, instead of a univariate time series embedded with some lag $\tau > 1$, we have a multidimensional time series with a lag $\tau = 1$, and rather than having a fixed threshold ϵ, we sweep through a range of values. The overall structure is similar: for a series X of length t, we construct a $t \times t$ binary matrix M where an element $M_{ij} = 1$ if $d(X(t_i), X(t_j)) < \epsilon$, which can be cast as a network [17]. Donner et al., (2010) detail how creating

networks from continuous dynamical systems can allow researchers to leverage the power of network science and graph theory to understand complex, nonlinear systems. We feel strongly that these sorts of analyses have many applications in neuroscience, biology, and other natural sciences.

It is worth discussing some of the limitations of this project. Arguably the most severe is that we have no way of knowing whether the monkeys in the ketamine condition were actually experiencing phenomenological consciousness. We are extrapolating from well-described reports from humans [33, 16, 66], but we recognize this as a fundamental limitation of this kind of animal-based consciousness research. The second limitation is that we only have a small number of unique individuals (two: Chibi and George) in each of the states (two ketamine scans, two propofol scans, four awake scans), from which we sample many slices. Given the rarity of this kind of dataset, a small sample size currently is unavoidable, however future work to generalize these findings may yield greater statistical power. Finally, for the OPN analysis, there is a fundamental question of how to select the optimal values for parameters d and τ [38, 41, 69]. While we have made an effort to make choices that are theoretically grounded and data-driven, further mathematical work may provide better insights and further unify the field. Finally, comparing ketamine and propofol anesthesia is somewhat low-hanging fruit in terms of dynamical-systems analysis of the brain: both NMDA and GABA receptors are widely expressed throughout CNS [23] and so altering activity at these receptors is expected to result in wide-scale changes to brain dynamics. Whether these methods are as effective when comparing more subtle differences in brain function, for example comparing clinically relevant states such as depression, OCD, or schizophrenia, remains to be seen.

Some aspects of this work will generalize to other neuroimaging modalities, while others will not. The EPC can be readily constructed from BOLD time series and the relative velocity analysis easily follows. The relative change in "velocity" induced by propofol and ketamine could be easily replicated in fMRI data, as could the persistence homology of the EPC analysis. In a similar vein, these analyses could also be adapted on EEG and MEG data. While constructing an OPN appears feasible for MEG and EEG data, its construction from BOLD time series presents problems, as OPNs generally require a long time series ($>> 1000$ samples) to capture a full representation of the state space.

5 Conclusions

Our understanding of how brain dynamics relate to consciousness can be informed by exploring multiple different representations of time series data (velocity, embedded point-clouds, ordinal partition networks, etc). In this study we show that normal consciousness, ketamine anesthesia, and propofol anesthesia differ significantly along multiple dimensions. By exploring these differences, we can find the domains in which different states are more similar to each-other and in which they are more diverse. We find that ketamine anesthesia maintains several significant similarities with normal waking awareness. Future studies may address possible links between these dynamic signatures and ketamine's tendency to induce dream-like states of "dissociative anesthesia" rather than the "true anesthesia" produced by propofol. These analyses represent a new branch of "network neuroscience", that leverages network tools to address problems related to brain dynamics and activity. We show

that state-based analyses in both the temporal and spatial domains can be informative about what the brain is doing and how its dynamics relate to awareness.

Author Contributions

AP and TFV conceptualized and designed the experiment. AP, TFV and VD performed data analysis. AP, TFV, and OS made figures, wrote and edited the paper.

Acknowledgements

TFV is supported by NSF-NRT grant 1735095, Interdisciplinary Training in Complex Networks and Systems. AP is supported by the Indiana University Network Science Institute (IUNI). OS is supported by NIH/NIMH grant 1RO1MH121978-01. We thank Joshua Faskowitz for his wit and wisdom throughout the process.

References

[1] Bandt, C. and Pompe, B. (2002). Permutation entropy: a natural complexity measure for time series. *Physical Review Letters*, 88(17):174102.

[2] Bubenik, P., Hull, M., Patel, D., and Whittle, B. (2019). Persistent homology detects curvature. *arXiv preprint arXiv:1905.13196*.

[3] Cao, Y., Tung, W.-w., Gao, J. B., Protopopescu, V. A., and Hively, L. M. (2004). Detecting dynamical changes in time series using the permutation entropy. *Physical Review E*, 70(4):046217. Publisher: American Physical Society.

[4] Carhart-Harris, R. L. (2018). The entropic brain - revisited. *Neuropharmacology*, 142:167–178.

[5] Carhart-Harris, R. L. and Friston, K. J. (2019). REBUS and the Anarchic Brain: Toward a Unified Model of the Brain Action of Psychedelics. *Pharmacological Reviews*, 71(3):316–344.

[6] Carhart-Harris, R. L., Leech, R., Hellyer, P. J., Shanahan, M., Feilding, A., Tagliazucchi, E., Chialvo, D. R., and Nutt, D. (2014). The entropic brain: a theory of conscious states informed by neuroimaging research with psychedelic drugs. *Frontiers in Human Neuroscience*, 8.

[7] Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308.

[8] Casali, A. G., Gosseries, O., Rosanova, M., Boly, M., Sarasso, S., Casali, K. R., Casarotto, S., Bruno, M.-A., Laureys, S., Tononi, G., and Massimini, M. (2013). A Theoretically Based Index of Consciousness Independent of Sensory Processing and Behavior. *Science Translational Medicine*, 5(198):198ra105–198ra105.

[9] Chazal, F. and Michel, B. (2017). An introduction to Topological Data Analysis: fundamental and practical aspects for data scientists. *arXiv:1710.04019 [cs, math, stat]*. arXiv: 1710.04019.

[10] Comolatti, R., Pigorini, A., Casarotto, S., Fecchio, M., Faria, G., Sarasso, S., Rosanova, M., Gosseries, O., Boly, M., Bodart, O., Ledoux, D., Brichant, J.-F., Nobili, L., Laureys, S., Tononi, G., Massimini, M., and Casali, A. G. (2019). A fast and general method to empirically estimate the complexity of brain responses to transcranial and intracranial stimulations. *Brain Stimulation*, 12(5):1280–1289.

[11] Conley, C. C. (1978). *Isolated invariant sets and the Morse index*. Number 38. American Mathematical Soc.

[12] Crutchfield, J. P. (1994). The calculi of emergence: computation, dynamics and induction. *Physica D: Nonlinear Phenomena*, 75(1):11–54.

[13] Crutchfield, J. P. (2012). Between order and chaos. *Nature Physics*, 8(1):17–24.

[14] Csardi, G. and Nepusz, T. (2006). The igraph software package for complex network research. *InterJournal*, Complex Systems:1695.

[15] Delorme, A. and Makeig, S. (2004). EEGLAB: an open source toolbox for analysis of single-trial EEG dynamics including independent component analysis. *Journal of Neuroscience Methods*, 134(1):9–21.

[16] Domino, E. F., Chodoff, P., and Corssen, G. (1965). Pharmacologic effects of CI-581, a new dissociative anesthetic, in man. *Clinical Pharmacology and Therapeutics*, 6:279–291.

[17] Donner, R. V., Zou, Y., Donges, J. F., Marwan, N., and Kurths, J. (2010). Recurrence networks—a novel paradigm for nonlinear time series analysis. *New Journal of Physics*, 12(3):033025.

[18] Eckmann, J.-P., Kamphorst, S. O., and Ruelle, D. (1987). Recurrence Plots of Dynamical Systems. *Europhysics Letters (EPL)*, 4(9):973–977.

[19] Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Goj, R., Jas, M., Brooks, T., Parkkonen, L., and Hämäläinen, M. (2013). MEG and EEG data analysis with MNE-Python. *Frontiers in Neuroscience*, 7.

[20] Gramfort, A., Luessi, M., Larson, E., Engemann, D. A., Strohmeier, D., Brodbeck, C., Parkkonen, L., and Hämäläinen, M. S. (2014). MNE software for processing MEG and EEG data. *NeuroImage*, 86:446–460.

[21] Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008). Exploring network structure, dynamics, and function using networkx. In Varoquaux, G., Vaught, T., and Millman, J., editors, *Proceedings of the 7th Python in Science Conference*, pages 11 – 15, Pasadena, CA USA.

[22] Hari, R. and Puce, A. (2017). *MEG-EEG Primer*. Oxford University Press.

[23] Hemmings, H. C., Riegelhaupt, P. M., Kelz, M. B., Solt, K., Eckenhoff, R. G., Orser, B. A., and Goldstein, P. A. (2019). Towards a Comprehensive Understanding of Anesthetic Mechanisms of Action: A Decade of Discovery. *Trends in Pharmacological Sciences*, 40(7):464–481.

[24] Hoel, E. and Hoel, E. P. (2017). When the Map Is Better Than the Territory. *Entropy*, 19(5):188.

[25] Hoel, E. P., Albantakis, L., and Tononi, G. (2013). Quantifying causal emergence shows that macro can beat micro. *Proceedings of the National Academy of Sciences*, 110(49):19790–19795.

[26] Hunter, J. D. (2007). Matplotlib: A 2d Graphics Environment. *Computing in Science Engineering*, 9(3):90–95.

[27] Kaczynski, T., Mischaikow, K., and Mrozek, M. (2006). *Computational Homology*. Applied Mathematical Sciences. Springer New York.

[28] Kennel, M. B. and Buhl, M. (2003). Estimating good discrete partitions from observed data: Symbolic false nearest neighbors. *Physical Review Letters*, 91(8):084102–084102.

[29] Khan, K. S., Hayes, I., and Buggy, D. J. (2014). Pharmacology of anaesthetic agents I: intravenous anaesthetic agents. *Continuing Education in Anaesthesia Critical Care & Pain*, 14(3):100–105.

[30] Klein, B. and Hoel, E. (2019). Uncertainty and causal emergence in complex networks. *arXiv:1907.03902 [physics]*. arXiv: 1907.03902.

[31] Krakovská, A., Mezeiová, K., and Budáčová, H. (2015). Use of false nearest neighbours for selecting variables and embedding parameters for state space reconstruction. *Journal of Complex Systems*, 2015.

[32] Kruskal, W. H. and Wallis, W. A. (1952). Use of Ranks in One-Criterion Variance Analysis. *Journal of the American Statistical Association*, 47(260):583–621.

[33] Krystal, J. H., Karper, L. P., Seibyl, J. P., Freeman, G. K., Delaney, R., Bremner, J. D., Heninger, G. R., Bowers, M. B., and Charney, D. S. (1994). Subanesthetic effects of the noncompetitive NMDA antagonist, ketamine, in humans. Psychotomimetic, perceptual, cognitive, and neuroendocrine responses. *Archives of General Psychiatry*, 51(3):199–214.

[34] Kumar, S. R., Kumar, A., Trikha, A., Anand, S. S., and Gantla, P. (2012). Higuchi fractal dimension as a measure of analgesia. *IJMEI*, 4:66–72.

[35] Lewis, L. D., Weiner, V. S., Mukamel, E. A., Donoghue, J. A., Eskandar, E. N., Madsen, J. R., Anderson, W. S., Hochberg, L. R., Cash, S. S., Brown, E. N., and Purdon, P. L. (2012). Rapid fragmentation of neuronal networks at the onset of propofol-induced unconsciousness. *Proceedings of the National Academy of Sciences of the United States of America*, 109(49):E3377–3386.

[36] Liu, Q., Ma, L., Fan, S.-Z., Abbod, M. F., and Shieh, J.-S. (2018). Sample entropy analysis for the estimating depth of anaesthesia through human EEG signal at different levels of unconsciousness during surgeries. *PeerJ*, 6.

[37] Mann, H. B. and Whitney, D. R. (1947). On a Test of Whether one of Two Random Variables is Stochastically Larger than the Other. *The Annals of Mathematical Statistics*, 18(1):50–60.

[38] McCullough, M., Small, M., Stemler, T., and Iu, H. H.-C. (2015). Time lagged ordinal partition networks for capturing dynamics of continuous dynamical systems. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 25(5):053101.

[39] Millman, K. J. and Aivazis, M. (2011). Python for Scientists and Engineers. *Computing in Science Engineering*, 13(2):9–12.

[40] Mittal, K. and Gupta, S. (2017). Topological characterization and early detection of bifurcations and chaos in complex systems using persistent homology. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, 27(5):051102.

[41] Myers, A., Munch, E., and Khasawneh, F. A. (2019). Persistent Homology of Complex Networks for Dynamic State Detection. *arXiv:1904.07403 [nlin, physics:physics]*. arXiv: 1904.07403.

[42] Nagasaka, Y., Shimoda, K., and Fujii, N. (2011). Multidimensional recording (MDR) and data sharing: an ecological open research and educational platform for neuroscience. *PloS One*, 6(7):e22561.

[43] Oliphant, T. E. (2007). Python for Scientific Computing. *Computing in Science Engineering*, 9(3):10–20.

[44] Patania, A., Petri, G., and Vaccarino, F. (2017). The shape of collaborations. *EPJ Data Science*, 6(1):1–16. Number: 1 Publisher: SpringerOpen.

[45] Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., Vanderplas, J., Passos, A., Cournapeau, D., Brucher, M., Perrot, M., and Duchesnay, E. (2011). Scikit-learn: Machine Learning in Python. *J. Mach. Learn. Res.*, 12:2825–2830.

[46] Perea, J. A. (2018). Topological Time Series Analysis. *arXiv:1812.05143 [cs, math]*. arXiv: 1812.05143.

[47] Perea, J. A. and Harer, J. (2015). Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of Computational Mathematics*, 15(3):799–838.

[48] Robins, V., Meiss, J. D., and Bradley, E. (2000). *Computational topology at multiple resolutions: foundations and applications to fractals and dynamics*. PhD thesis, University of Colorado.

[49] Rosas, F., Mediano, P. A. M., Ugarte, M., and Jensen, H. J. (2018). An Information-Theoretic Approach to Self-Organisation: Emergence of Complex Interdependencies in Coupled Dynamical Systems. *Entropy*, 20(10):793.

[50] Rosvall, M. and Bergstrom, C. T. (2008). Maps of random walks on complex networks reveal community structure. *Proceedings of the National Academy of Sciences*, 105(4):1118–1123.

[51] Ruiz de Miras, J., Soler, F., Iglesias-Parro, S., Ibáñez-Molina, A. J., Casali, A. G., Laureys, S., Massimini, M., Esteban, F. J., Navas, J., and Langa, J. A. (2019). Fractal dimension analysis of states of consciousness and unconsciousness using transcranial magnetic stimulation. *Computer Methods and Programs in Biomedicine*, 175:129–137.

[52] Sarasso, S., Boly, M., Napolitani, M., Gosseries, O., Charland-Verville, V., Casarotto, S., Rosanova, M., Casali, A. G., Brichant, J.-F., Boveroux, P., Rex, S., Tononi, G., Laureys, S., and Massimini, M. (2015). Consciousness and Complexity during Unresponsiveness Induced by Propofol, Xenon, and Ketamine. *Current Biology*, 25(23):3099–3105.

[53] Saul, N. and Tralie, C. (2019). Scikit-tda: Topological data analysis for python.

[54] Schartner, M., Seth, A., Noirhomme, Q., Boly, M., Bruno, M.-A., Laureys, S., and Barrett, A. (2015). Complexity of Multi-Dimensional Spontaneous EEG Decreases during Propofol Induced General Anaesthesia. *PLOS ONE*, 10(8):e0133532.

[55] Schartner, M. M., Carhart-Harris, R. L., Barrett, A. B., Seth, A. K., and Muthukumaraswamy, S. D. (2017). Increased spontaneous MEG signal diversity for psychoactive doses of ketamine, LSD and psilocybin. *Scientific Reports*, 7:46421.

[56] Sizemore, A. E., Phillips-Cremins, J., Ghrist, R., and Bassett, D. S. (2018). The importance of the whole: topological data analysis for the network neuroscientist. *arXiv:1806.05167 [q-bio]*. arXiv: 1806.05167.

[57] Small, M. (2013). Complex networks from time series: Capturing dynamics. In *2013 IEEE International Symposium on Circuits and Systems (ISCAS2013)*, pages 2509–2512.

[58] Traile, C., Saul, N., and Bar-On, R. (2018). Ripser.py: A Lean Persistent Homology Library for Python. *The Journal of Open Source Software*, 3(9):295.

[59] Tsuchiya, N., Taguchi, S., and Saigo, H. (2016). Using category theory to assess the relationship between consciousness and integrated information theory. *Neuroscience Research*, 107:1–7.

[60] Varley, T. F., Craig, M., Adapa, R., Finoia, P., Williams, G., Allanson, J., Pickard, J., Menon, D. K., and Stamatakis, E. A. (2020a). Fractal dimension of cortical functional connectivity networks & severity of disorders of consciousness. *PLOS ONE*, 15(2):e0223812. Publisher: Public Library of Science.

[61] Varley, T. F., Luppi, A. I., Pappas, I., Naci, L., Adapa, R., Owen, A. M., Menon, D. K., and Stamatakis, E. A. (2020b). Consciousness & Brain Functional Complexity in Propofol Anaesthesia. *Scientific Reports*, 10(1):1–13. Number: 1 Publisher: Nature Publishing Group.

[62] Walt, S. v. d., Colbert, S. C., and Varoquaux, G. (2011). The NumPy Array: A Structure for Efficient Numerical Computation. *Computing in Science Engineering*, 13(2):22–30.

[63] Wang, J., Noh, G.-J., Choi, B.-M., Ku, S.-W., Joo, P., Jung, W.-S., Kim, S., and Lee, H. (2017). Suppressed neural complexity during ketamine- and propofol-induced unconsciousness. *Neuroscience Letters*, 653:320–325.

[64] Wassermann, G., Arnol'd, V., Dzhamay, A., Novikov, S., Dubrovin, B., Givental', A., Kirillov, A., and Krichever, I. (2001). *Dynamical Systems IV: Symplectic Geometry and its Applications*. Encyclopaedia of Mathematical Sciences. Springer Berlin Heidelberg.

[65] Yanagawa, T., Chao, Z. C., Hasegawa, N., and Fujii, N. (2013). Large-Scale Information Flow in Conscious and Unconscious States: an ECoG Study in Monkeys. *PLoS ONE*, 8(11).

[66] Zanos, P., Moaddel, R., Morris, P. J., Riggs, L. M., Highland, J. N., Georgiou, P., Pereira, E. F. R., Albuquerque, E. X., Thomas, C. J., Zarate, C. A., and Gould, T. D. (2018). Ketamine and Ketamine Metabolite Pharmacology: Insights into Therapeutic Mechanisms. *Pharmacological Reviews*, 70(3):621–660.

[67] Zhang, J., Zhou, J., Tang, M., Guo, H., Small, M., and Zou, Y. (2017). Constructing ordinal partition transition networks from multivariate time series. *Scientific Reports*, 7(1):7795.

[68] Zhang, X. S., Roy, R. J., and Jensen, E. W. (2001). EEG complexity as a measure of depth of anesthesia for patients. *IEEE transactions on bio-medical engineering*, 48(12):1424–1433.

[69] Zou, Y., Donner, R. V., Marwan, N., Donges, J. F., and Kurths, J. (2019). Complex network approaches to nonlinear time series analysis. *Physics Reports*, 787:1–97.