

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

Thomas F. Varley^{1,2}, Vanessa Denny¹, Olaf Sporns^{1,3}, Alice Patania³

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 electrophysiological activity recorded from cerebral cortex of two non-human primates. 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 that the awake condition generally had the "richest" structure, with the widest repertoire of available states, the presence of pronounced higher-order structures, and the least deterministic dynamics. In contrast, the propofol condition had the most dissimilar dynamics to normal consciousness, transitioning to a more impoverished, constrained, low-structure regime. The ketamine condition, interestingly, seemed to combine aspects of both: while it was generally less complex than the awake condition, it remained well above propofol in almost all measures. These results may provide insights into how consciousness can persist under the influence of ketamine and the battery of measures used provides deeper and more comprehensive insights than what is typically gained by using point-measures of complexity.

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 [61]. Pharmacology has revealed the molecular actions of different drugs [31], 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 [54, 64]. 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, they are useful to compare due to their markedly different pharmacologies and the differences between the states they induce at low-to-moderate doses. Propofol is one of the most commonly-used anesthetics in medicine, and while its full mechanism of action is not totally understood, a key feature is believed to be its widespread modulation of GABA_A receptors [31, 24]. By binding to GABA_{Ar}s, propofol potentiates the effects of endogenous GABA, causing widespread inhibition of neuronal activity. Consequently, even at low doses, propofol induces states of amnesia, sedation, atonia, and at higher doses, full anesthesia. Propofol reliably induces a suite of changes to oscillatory activity in the brain, such as the emergence of widespread frontal alpha patterns [1], as well as regionally specific changes to gamma and beta bands indicative of a multi-stage fragmenting of neural communication networks [29].

In contrast to propofol, ketamine acts primarily as an antagonist of glutamatergic NMDA receptors, which has a local excitatory response [31, 67]. Blockade of NMDA receptors has been found to disinhibiting activity in cortical pyramidal neurons, driving activity in local cortical circuits [27]. Ketamine causes widespread, light central nervous system stimulation, in contrast to propofol's deeply sedating properties. Despite this increase in activity, ketamine has been found to significantly disrupt directed information flow across the cortex [55] in a manner consistent with the network fragmentation observed under GABAergic surgical anaesthetics [35, 36]. Like propofol, ketamine is associated with a distinct pattern of changes to oscillatory power with a particular focus on the emergence of high frequency activity in the 20-70Hz band, which distinguishes it from other commonly used surgical anaesthetics [38] and suggests that the different pharmacologies of the two drugs result in different "paths to unconsciousness," described by different dynamical regimes [2]. The state that ketamine induces is typically referred to as "dissociative anesthesia" and represents a highly atypical state of consciousness [16, 34]. In a state of dissociative anesthesia, an individual will often appear to be unresponsive to stimuli (including pain) and, to an external observer, may be indistinguishable from someone anesthetized with a typical anesthetic like propofol. 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 [67].

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 [64]. 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 [53, 51]. In contrast, ketamine’s blockade of NMDA receptors disinhibits cortical neurons, causing widespread, uncoordinated excitatory activity [27, 67, 54, 64]. 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 [8, 7] and received empirical support from studies of psychedelic drugs [6]. The majority of these studies rely ultimately on point-summaries of “complexity” (eg: Lempel-Ziv complexity [53, 51], entropy [37, 69], 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 [44, 66]. 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 , 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. In addition to these local measures, modeling continuous brain dynamics 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 (TDA) [10, 56] 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, 49, 65, 48] 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, 40, 43] (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 [44]. Specifically, we used data from two monkeys (Chibi and George), prior to, and during behavioural unresponsiveness with propofol or ketamine [66]. 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 [66], and can be viewed on the relevant Neurotycho wiki page¹. 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, 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, artifact-free ten-second intervals from the anesthetized condition and three artifact-free ten-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). Ten-second intervals were the largest that we found computationally tractable: the run-time and memory requirements of the topological data analysis grows super-exponentially with the number of samples, and consequently this required us to keep them comparatively short. 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].

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

2.1.2 Pre-Processing

After subsections of the time series had been selected, pre-processing was performed in MNE-Python [20, 21]. Each series was high-pass filtered with a low frequency of 0.1 Hz, 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 [23]. 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 (in this case, 9900). 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, C -dimensional state space. For visualization of this process, see Figure 1. 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 [9, 46, 48]. 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

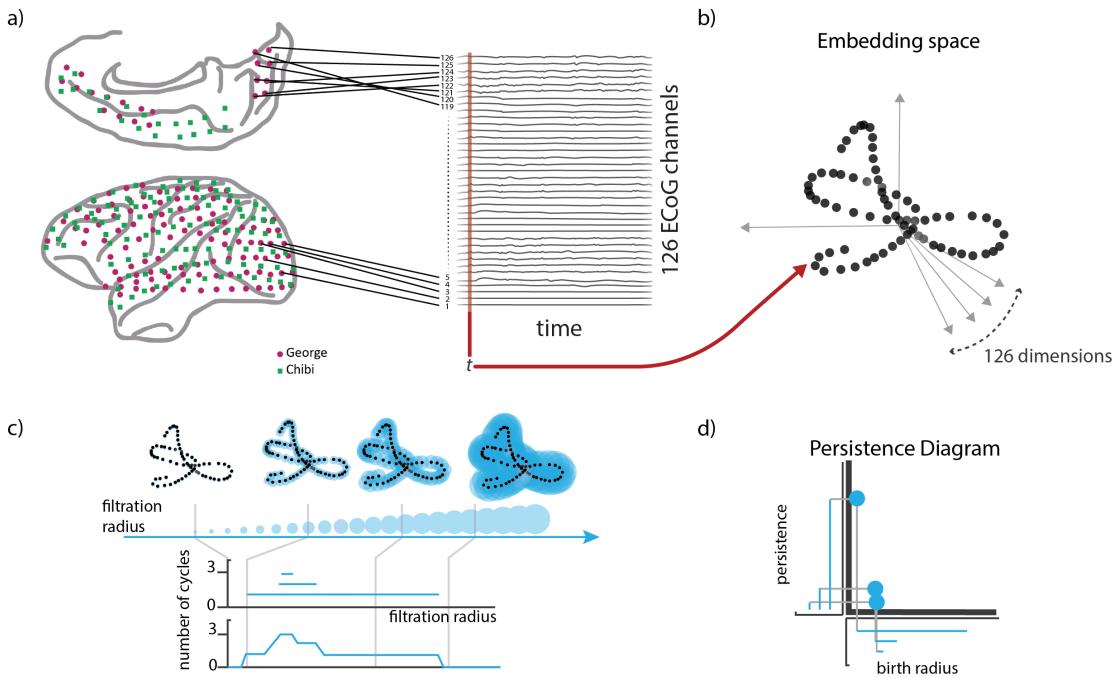


Figure 1: Construction of the embedded point cloud from a multidimensional time series. Every time-step in the recording corresponds to a single column vector in a $C \times T$ array (a), and this vector can be imagined as a point in a C -dimensional space, creating a point-cloud that traces out the 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. 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. To classify the evolving manifold we use two summaries: the Betti curve, that counts the number of cycles present in the discretized manifold for each filtration radius (c); and the persistence diagram, that encodes each cycle as a point and records the radius of its first appearance (birth radius) in the x axis and until which radius it persists in the y axis (d).

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.

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 [4, 5]. 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 [4].

To explore the temporal dynamics of the system, we constructed ordinal partition network (OPN) representations of the data [57, 40], following the procedure discussed in [43]. Due to the difficulties associated with multi-dimensional OPNs [68], we applied this method to every channel in every recording slice individually, and then aggregated the results. This allows us to capture all of the information in each single channel, as multi-channel analysis with OPNs is not technically feasible.

To construct an OPN, begin with a time series $X = x_1, x_2, \dots, x_n$. 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 $v_i = [x_i, x_{i+\tau}, \dots, x_{i+(d-1)\tau}]$. 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 $v_{i,1} \leq v_{i,2} \leq \dots \leq v_{i,d}$. 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 $1, 2, \dots, d$ $n_i = [\pi(x_i), \pi(x_{i+\tau}), \dots, \pi(x_{i+(d-1)\tau})]$. 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. 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 [40], false-nearest-neighbours criteria [30], or the first local maxima in the permutation entropy [43]. 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

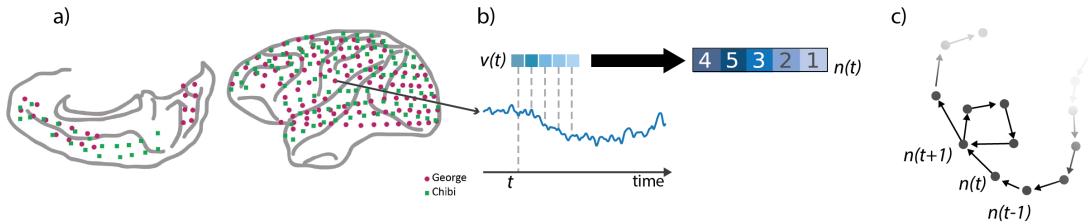


Figure 2: We constructed the OPN of the projection of each 2s interval multi-variate time series onto the first temporal principal component of the entire data set (a). An OPN was computed for each channel separately. Each channel time series is embedded in a d dimensional space, using a time-lag τ , as is done when constructing Taken’s Embedding into the phase space. 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 $1, 2, \dots, d$. 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).

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.

As OPNs are largely restricted to one-dimensional time-series, we created a unique OPN for each channel in each of the scans. To account for natural differences in the dynamics of each channel or brain region, each OPN was constructed with a unique embedding lag τ . There were no significant differences between any conditions at the subject level in terms of the average optimal lag. In contrast, every OPN had to be constructed with the same embedding dimension d to enable proper comparison and data aggregation. We selected τ using the first zero crossing of the autocorrelation, and $d = 5$ was the mode of the distribution of embedding dimensions for which the variance in the degree distribution was maximal. By constructing networks at the channel level, we can assess changes in brain dynamics at a number of scales, including at the channel-level, the region level, and at the individual level by aggregating different numbers of channels. In the Results section, the summary statistics are aggregated at the subject level unless otherwise specified.

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 two measures of how much information a system encodes in its state transition graph [26, 25, 32]. 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:

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

These two measures constitute an information-theoretic analysis of state-transition graphs representing a system outputting a continuous signal. The OPN can be thought of as an approximation 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 [26, 32]. 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) [45, 42]. Analysis of variance was computed using the Kruskal-Wallis Analysis of Variance test [33], and post hoc testing was done using the Mann-Whitney U test [39]. Non-parametric statistics were chosen due to the uncertainty that the data were sufficiently Gaussian. OPNs were constructed using the OPyN package (available on Github: <https://github.com/thosvarleyf/OPyN>). Persistence homology analysis was done using the Ripser (version 0.3.2) [60] and Persim libraries (version 0.0.9) as part of the Scikit-TDA library (version 0.0.4) [52]. Other packages used include the Numpy library (version 1.15.4) [63], Scikit-Learn (version 0.20.0) [47], Matplotlib (version 2.2.2), [28], Spyder (version 3.2.3), NetworkX (version 2.2) [22], 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

To characterize how propofol and ketamine altered brain dynamics in spacial and temporal domains, we constructed two embedded representations of the data. The first was the embedded point cloud (EPC), which embeds the instantaneous activity across all channels in a high-dimensional space as a point-cloud, which can then be analyzed using techniques from topological data analysis. The second method, which focuses on temporally extended dynamics, involves constructing a representative state-transition network for each channel (ordinal partition networks, OPNs), which encodes temporally extended dynamics in its structure. We begin by discussing the results of the EPC analysis, and then move on to the OPNs.

Measure	Category	Formalism
Maximum Persistence	EPC/TDA	The length of the longest lived cycle in the Rips filtration
Total Number of Cycles	EPC/TDA	The total number of cycles that appears over the Rips filtration
Maximum Number of Cycles	EPC/TDA	The maximum number of cycles present at a given moment
Number of Nodes	OPN	The number of unique nodes in the OPN.
Number of Edges	OPN	The number of unique edges in the OPN.
Determinism	OPN/Info. Theory	The average entropy of the out-going edges for each node.
Degeneracy	OPN/Info. Theory	How much information is lost when states run together.
Modularity	OPN	How well the nodes in the OPN can be clustered.
Permutation Entropy	Timeseries/Info. Theory	The entropy of the permutation-embedded series.
Lyapunov Exponent	Timeseries	The "chaoticity" of the time-series.

Table 1: A table describing all of the measures described here, and how they can be intuitively interpreted. The measures are broadly categorized into several groups, including which embedding they are applied to (EPC vs. OPN), and the general mathematical fields they are derived from (EPC, Information Theory)

Condition	Maximum Persistence	Number of Cycles	Maximum Number of Cycles
Awake	0.36 ± 0.09	417.5 ± 119.61	80.0 ± 29.36
Ketamine	0.4 ± 0.11	190.88 ± 39.24	22.46 ± 8.09
Propofol	0.4 ± 0.11	64.17 ± 19.47	10.125 ± 2.8

Table 2: Results for the three measures used on the EPC: the maximum persistence of the longest lived cycle, the total number of cycles over the course of the whole filtration, and the maximum number of cycles to exist at any one point in the filtration.

3.1 Embedded Point Cloud

Kruskal-Wallis analysis of variance found no significant differences between the three conditions for the maximum persistence of the longest cycle.

There were significant differences between all the conditions regarding the total number of cycles to exist over the course of the filtration ($H=62.8$, $p=2.3 \times 10^{-14}$). The largest total number of cycles was found in the awake condition (417.5 ± 119.61), followed by the ketamine condition (190.88 ± 39.24), and the propofol condition had the fewest number of cycles (64.17 ± 19.47). A similar pattern held for the maximum number of cycles to exist at any individual point in the filtration: Kruskal-Wallis analysis of variance found significant differences between all three conditions ($H=59.46$, $p=1.22 \times 10^{-13}$), with the awake condition having the most cycles (80.0 ± 29.36), followed by the ketamine condition (22.46 ± 8.09), and with the propofol condition having the fewest (10.125 ± 2.8). These results suggest that the propofol condition has the least amount of "structure" constraining the simultaneous evolution of activity across the channels. Recall that, if every channel were acting independently, the resulting EPC would be a smooth, multivariate Gaussian distribution in as many dimensions as there are channels. The existence of cycles suggests a deviation from this maximally entropic ideal - the evolution of the channels appear to be jointly constrained by each-other, creating cycles and voids. All results are recorded in Table 2 and visualized in Figure 4.

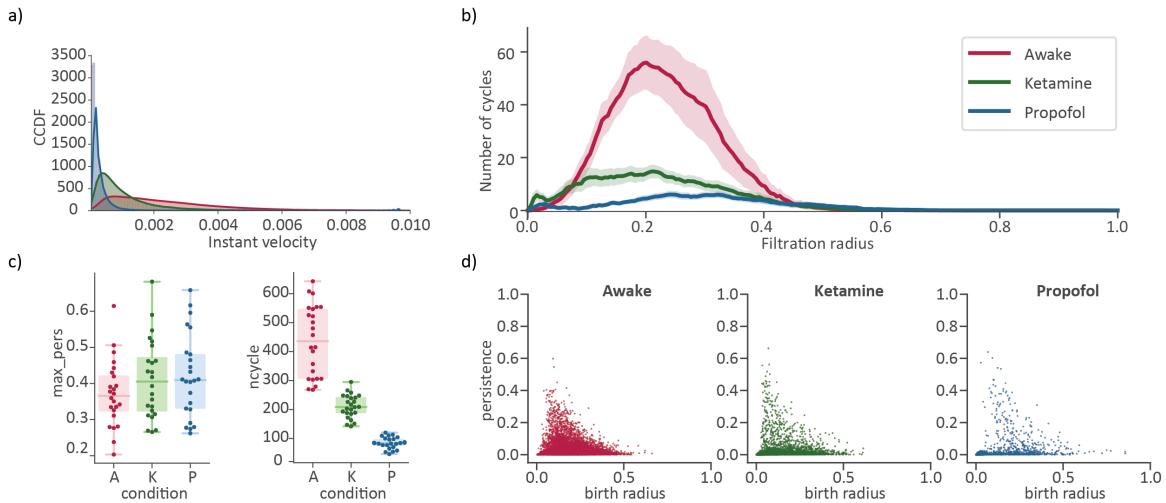


Figure 3: Results for the EPC analysis of brain dynamical trajectories in the high-dimension configuration space. Empirical distributions of the instantaneous velocity for the three conditions (a). Avarage Betti curves, with a 90% confidence interval, counting the number of homological cycles present in the complex as the spheres radius increases (b). The life duration of the longest cycle in each time segment and total number of homological cycles (represented by a point in the swarm plots), separated by condition (A = Awake, K = Ketamine, P = Propofol)(c). Persistence diagrams for all time series embedding separated by conditions. For each condition a scatter point represents the birth radius and persistence of a cycle in one of the embedding point cloud(d).

Condition	Number of Nodes	Number of Edges	Determinism	Degeneracy	Modularity
Awake	102.28 ± 28.9	480.67 ± 208.95	0.91 ± 0.03	0.0112 ± 0.012	0.81 ± 0.08
Ketamine	86.41 ± 30.14	325.13 ± 155.86	0.86 ± 0.035	0.02 ± 0.01	0.84 ± 0.08
Propofol	65.17 ± 26.6	191.15 ± 95.46	0.82 ± 0.04	0.0114 ± 0.009	0.85 ± 0.13

Table 3: Results for the five measures used to characterize the OPNs: the number of nodes in the network, the number of edges, the determinism, degeneracy, and modularity of the network. Each of these measures can be thought of as a different axes along which the discrete state-transition dynamics can occur.

3.2 Ordinal Partition Network

We assessed five simple network measures to characterize how the topology of the OPNs changed between conditions. The simplest measure is the number of nodes, which measures the size of the repertoire of ordinal partition micro-states available to the system over the course of it's run. Kruskal-Wallis analysis of variance found significant differences between all three conditions ($H=2296.06$, $p < 10^{-20}$), with the awake condition having the most nodes (102.28 ± 28.9), followed by the ketamine condition (86.41 ± 30.14), with the propofol condition having the fewest (65.17 ± 26.6). This is consistent with the original Entropic Brain Hypothesis, that the vividness of consciousness, and complexity of behavior, tracks the size of the repertoire of available states [8, 6].

We also compared the number of edges present in the network. In the same way that the number of nodes counts the unique micro-states the system adopts, the number of edges counts

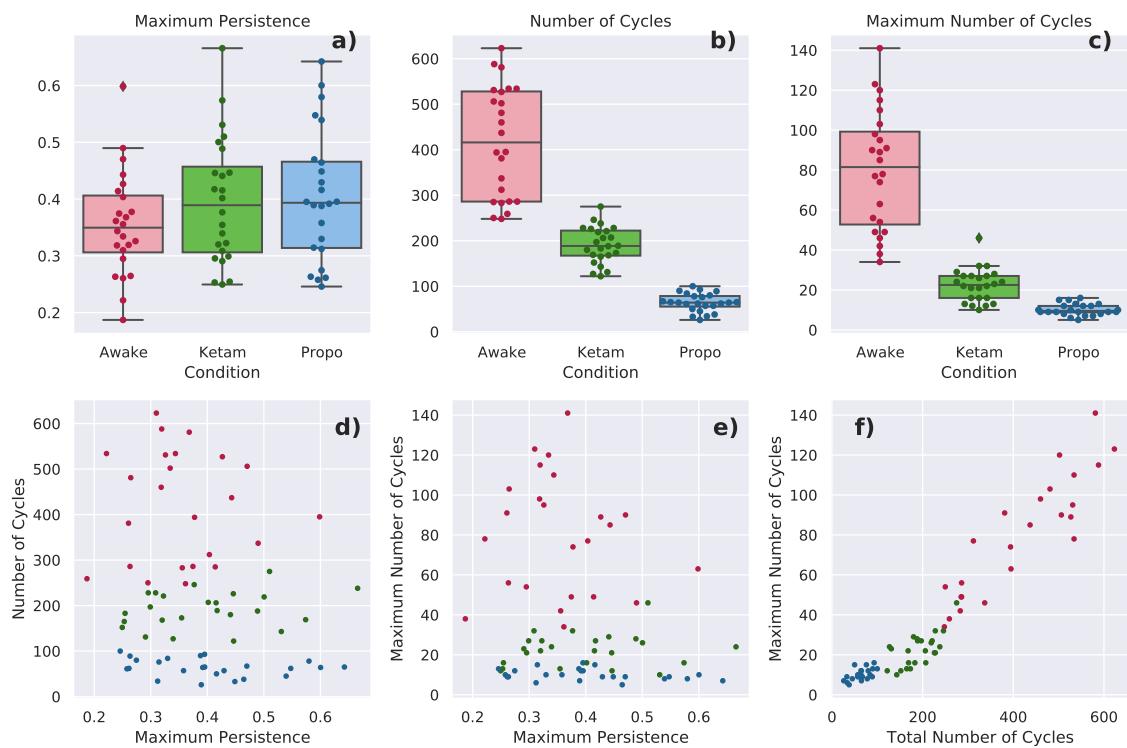


Figure 4: a-b) Boxplots showing how the three conditions differed on each of the TDA analyses applied to the EPC: the maximum persistent (a), the total number of cycles (b), and the maximum number of cycles (c). d-f) scatter plots showing how the three measures above relate to each-other. Note the strong positive relationship in f), indicating that the total number of cycles is related to the largest number of cycles that appear at a given moment. "Cycles" in this context can be understood as constraints on the behavior of the system indicating a deviation from the case where all channels are acting independently.

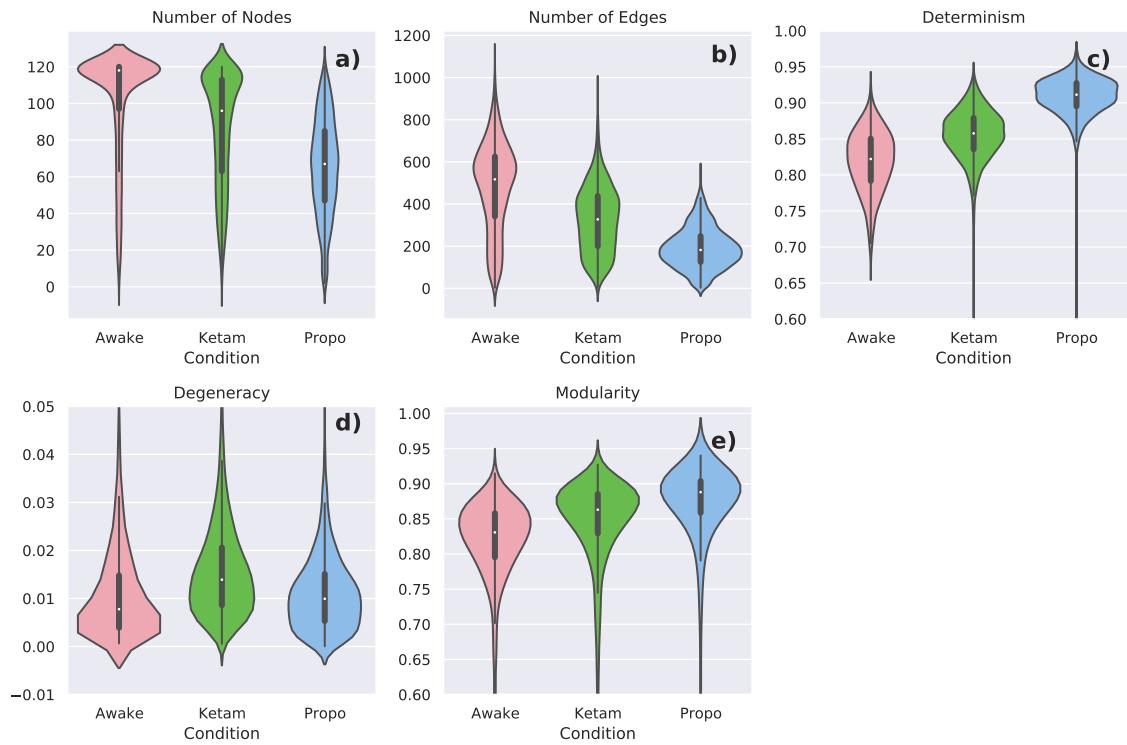


Figure 5: Violin plots for the various OPN measures. Violin plots were chosen to represent the distribution of measures over thousands of individual channels. a) The difference between the three conditions in the number of unique nodes present in the OPN. This represents something like a measure of the size of the repertoire of states available to the system in a given condition. b) The number of unique edges in the OPN, represents the flexibility of the systems dynamics. c) The determinism: an information-theoretic measure of how reliably the future of the system can be predicted from its' past. d) the degeneracy: a measure of how much information about the past is lost when different states "run into" each-other. This is a rare case where the ketamine condition is higher than either of the awake or propofol conditions. e) the modularity of the OPN network, determined using the Infomap algorithm [50]. Provides a measure of how "constrained" the systems dynamics are by higher-order patterns of state-transitions.

Condition	Permutation Entropy	Lyapunov Exponent
Awake	4.27 ± 0.59	0.2 ± 0.013
Ketamine	3.92 ± 0.58	0.18 ± 0.01
Propofol	3.61 ± 0.61	0.16 ± 0.01

Table 4: Results for the two "classical" measures used to characterize the chaoticity and information content in timeseries: the permutation entropy and the Lyapunov exponent.

the unique transitions that the system can perform. There were significant differences between all three conditions ($H=2645.41$, $p < 10^{-10}$, with the awake condition having the most unique edges (480.67 ± 208.95), followed by the ketamine condition (325.13 ± 155.86) and then the propofol condition (191.15 ± 95.46). This suggests that, in addition to the larger repertoire of individual states, there is also increased flexibility in terms of how those states transition between themselves.

To leverage the natural application of information theoretic analysis to OPNs, we used two measures originally developed to assess the causal structure of the system: determinism (on average, how predictable is the future from the present) and degeneracy (on average, how well can the past be reconstructed from the present) [32]. There were significant differences between all three conditions ($H=4678.4$, $p < 10^{-20}$, however in contrast to previously described measures, it was the propofol condition that had the highest determinism (0.82 ± 0.04 bit), followed by the ketamine condition (0.86 ± 0.035 bit), and the awake condition had the lowest determinism (0.91 ± 0.03 bit). The degeneracy was also significantly different between conditions ($H=607.94$, $p < 10^{-20}$), with ketamine having the highest degeneracy (0.02 ± 0.01 bit), followed by propofol (0.0114 ± 0.009 bit), and then the awake condition (0.0112 ± 0.012 bit).

The final network-measure we applied was the modularity, using the Infomap modularity algorithm [50]. The Infomap algorithm assigns a subset of nodes to the same community if a random walker on the network has a tendency to get "stuck" in that subset - in the context of a state-transition network, where a random walk is naturally understood as a possible trajectory of the system through state-space, a module could then be understood as a kind of "metastable attractor" that the system gets transiently caught in. A high modularity, then, is indicative of strong higher-order attractor dynamics constraining the evolution of the system, while a low modularity describes a relatively "flat" state transition landscape. Kruskal-Wallis analysis of variance found significant differences between all three conditions ($H=1734.71$, $p < 10^{-20}$), with the propofol condition having the highest modularity (0.85 ± 0.13 , followed by the ketamine condition (0.84 ± 0.08), and then the awake condition had the least modular structure (0.81 ± 0.08). All results for this section are tabulated in Table 3 and visualized in 5.

3.2.1 Raw Timeseries Measures

In addition to the analysis of the ordinal partition networks themselves, we performed two classical non-network based analyses of the timeseries, to compare how our novel methods compared to more established ones. The first was the permutation entropy [4] (which is intimately related to the construction of the OPN). Kruskal-Wallis analysis of variance found significant differences between

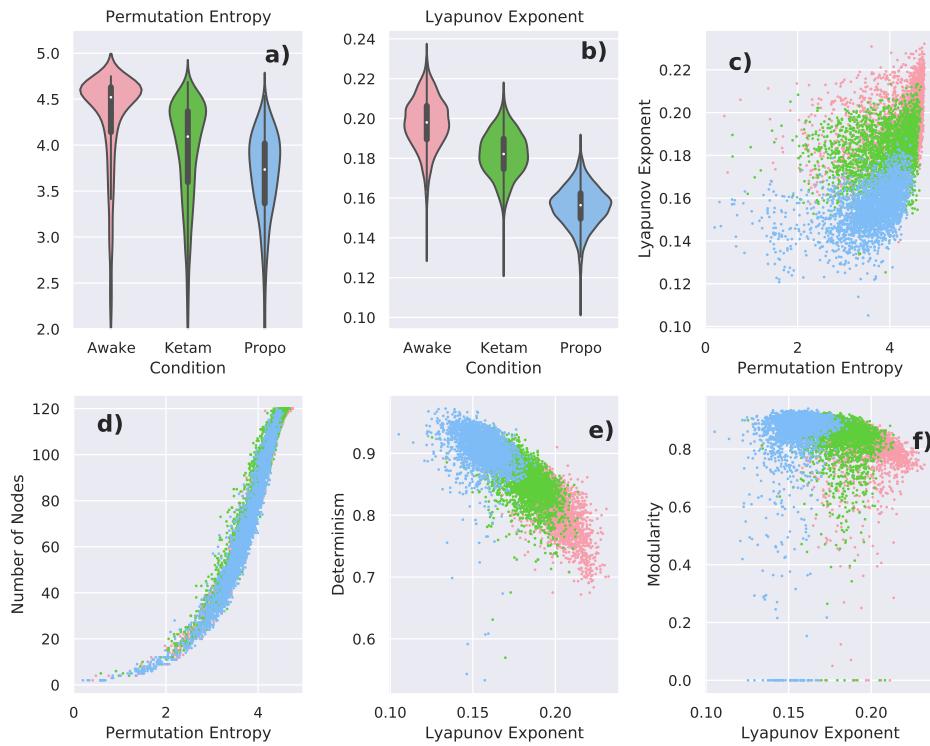


Figure 6: a-b): Violin plots of the permutation entropy and the Lyapunov exponent. As with the OPN measures, we chose violin plots to capture the distribution of a large number of channels. C) a scatter plot of the permutation entropy against the Lyapunov exponent: note the positive, but non-linear relationship between both measures, consistent with previous work showing an association between permutation entropy and chaos [4]. d-f): scatter plots showing how the two time-series measures relate to the OPN measures. The scatter-plots show that that these novel measures derived from the OPN are consistent with established measures. As expected, the number of nodes in positively associated with the permutation entropy (d). The Lyapunov exponent is negative correlated with the determinism (which is consistent with the results reported in [62] and consistent with the intuition behind chaotic systems (e). Finally, there is a slight negative relationship between the Lyapunov exponent and the modularity, suggesting that metastable higher-order dynamics may be harder to maintain in chaotic systems (f).

all three conditions (statistic=2219.65, $p < 10^{-20}$). As usual, the awake condition had the highest permutation entropy (4.27 ± 0.59 bit), followed by the ketamine condition (3.92 ± 0.58 bit), and with propofol having the lowest (3.61 ± 0.61). These results indicate that, on average, the system is visiting all of the realized microstates more equitably in the awake condition, whereas in the propofol condition, it is repeatedly returning to a subset of the states and only rarely visiting others.

The other measure we used was the Lyapunov exponent, commonly understood as a measure of “how chaotic” a system is. Once again, the usual pattern held: significant differences between all three conditions ($H=5628.26$, $p < 10^{-20}$), with the awake having the greatest chaoticity (0.2 ± 0.013), following by the ketamine condition (0.18 ± 0.01), and the propofol condition (0.16 ± 0.01). Chaoticity can be thought of as something like a measure of how sensitive a system is to perturbation (how rapidly a perturbed trajectory diverges from its unperturbed self). This may be naturally understood in the context of the need for conscious, awake, animals to be able to rapidly respond to new stimuli from the environment. Sensitivity to environmental perturbations has clear benefits, although in the case of ‘overly-chaotic dynamics’, it would certainly become detrimental. All results are tabulated in Table 4 and visualized in Figure 6.

3.3 Dimensionality Reduction & Visualization

We can imagine that each one of the results discussed above defines a kind of ”dynamical morphospace”, analogous to previous work that’s been done on morphospaces in network topology [3], where every network is embedded in a high-dimensional space: the value along a given axis is defined by the various measures described above (i.e. one axis is determinism, one is degeneracy, etc). In the case of just the OPN results, this gives us one point for every channel embedded in a 5-dimensional morphospace. This can be visualized using a dimensionality reduction algorithm like PCA, tSNE, or UMAP to create a “birds-eye-view” of how the different states of consciousness relate to each-other. While this does not return a quantitative measure of similarity or difference, it provides a useful visualization of how the different conditions are related to each-other, which can be more intuitive than a table of numbers. We used the UMAP embedding algorithm [41] to construct a visualization of the channel-wise relationships between the conditions using just the OPN data (see Fig. 7).

4 Discussion

In this paper, we have discussed several ways in which temporal and spatial, embedding of electrophysiological data from macaques in three distinct states of consciousness (awake, ketamine anesthesia, propofol anaesthesia) can reveal insights into how brain dynamics reflect alterations to consciousness along many axes. Historically, there has been considerable interest in one-dimensional, scalar measures of how the “complexity” of brain activity relates to consciousness (e.g. Lampel-Ziv complexity [53, 54] or Integrated Information Theory’s “Phi” value [59]). However, as argued by Feldman and Crutchfield [18], there are fundamental limitations to how much insight can be gained by even a “well-behaved” scalar measure of complexity. When attempting to characterize a system as complex as a conscious (or even unconscious) brain which can vary along many different axes, a more compe-

hensive picture is necessary. By constructing two embeddings (the EPC and the channel-wise OPNs) we can paint a much more holistic picture of how brain activity changes in spatial and temporal domains when consciousness is lost or altered. Rather than creating a ranking of “complexity” from low to high, we can begin to tease out the ways in which these conditions are similar, and different.

The awake condition could be characterized as having a high degree of interaction between the individual channels when compared to propofol, as evidence by the persistence homology results: the presence of a large number of cycles suggests that the dynamics of the individual regions are subject to collective constraints that are less prominent when consciousness is lost. At the level of individual channels, the awake condition can adopt the largest number of unique micro-states and has the highest degree of flexibility transitioning between them. At the macro-scale, modularity analysis revealed that the state-transition network has the lowest modular structure, which indicates that the system is less likely to get caught in deep attractors compared to the propofol condition. Finally, the awake condition is both less deterministic and less degenerate, which previous work has found to be indicative of the onset of chaotic dynamics [62], an interpretation supported by the Lyapunov exponents as well. This suggests that the temporal dynamics of the awake condition are the least predictable, suggesting a high-degree of flexibility compared to either anesthesia states.

We can compare these findings with the results from the propofol analysis. Persistence homology analysis of the EPC found that the propofol condition had the lowest number of cycles, suggesting a loss of “higher-order structure” driving activity across multiple channels. This is consistent with previous findings that propofol anaesthesia decreases functional connectivity [19], however, one significant benefit of the persistence homology analysis is that it considers the joint-states of all channels together, as opposed to examining pairwise relationships between individual channels, which may miss higher-order synergies that may be present in the system [58, 17]. The propofol condition had the smallest repertoire of available states and comparatively constrained transitions between them. Modularity analysis bore this out, finding that the propofol condition had a significantly higher modularity, suggesting that the system is more likely to get “stuck” in subsets of the state space. It was also the most deterministic suggesting reduced dynamical flexibility. Interestingly, it was the ketamine condition that had the most degenerate dynamics, suggesting that the ketamine has the shortest ”memory”, as the past states are minimally predictable from the present.

In general, the ketamine condition occupied something of a middle ground between the awake and propofol conditions, suggesting that it combines elements of both states in its dynamics. This is clearly visible when the UMAP embedding is performed on the OPN-morphospace: the ketamine condition is clearly visible forming a kind of boundary between the awake and propofol conditions, which do not overlap significantly. This is consistent with the known clinical properties of ketamine anaesthesia: while it produces a state that is externally very similar to propofol anaesthesia (loss of responsiveness to stimuli, analgesia, etc), ketamine can produce dream-like, dissociative states [34, 31], suggesting that the process generating phenomenological consciousness is not completely inhibited.

This work does have limitations which are worth considering. The most obvious is the small sample size: two macaques is a small N , even with multiple slices taken out of the longer scans. Given the origin of this data, this limitation cannot be currently addressed and we hope that these

results, and the larger methods introduced, can replicate these findings in future studies. We also cannot directly infer what state of consciousness the macaques were in at any given time, or even if macaques are capable of experiencing something like the dissociative anaesthesia that ketamine induces in humans. As with the small N , this is something of a fundamental limitation and an on-going issue in consciousness research. In terms of the OPNs, a significant limitation is that they can only be constructed from a single channel: while attempts at multivariate generalizations have been proposed [68] 128-channel systems such as those explored here remain computationally and practically prohibitive. This highlights the importance of multiple different measures to bear on a question, as opposed to looking for a singular test that explains “everything.” While the OPNs are limited to one channel, when combined with the EPC analyses, we get a much richer picture combining temporal and spatial dynamics into a single whole.

We anticipate that the techniques described here can be used to understand other states of consciousness, such as psychedelia or disorders of consciousness following brain injury, as well as understanding individual differences in normal cognition. We might hypothesize, for instance, that high performance on creativity tasks might be associated with an increase in the repertoire of microstates discernible by the OPN and the flexibility with which the brain transitions between them. Using the notion of a dynamical morphospace, it may be possible to create a “map” of different cognitive processes based on their dynamical similarities and differences.

5 Conclusions

In this work, we describe how embeddings of neural activity data can help differentiate between the similarities and differences between three distinct states of consciousness: normal waking awareness, propofol anaesthesia, and ketamine anesthesia. To assess the spatial distribution of activity across channels, we used topological data analysis to analyze the structure of the joint-states of all channels through time. To assess the channel-level temporal dynamics, we construct discrete state-transition graphs using ordinal partition networks, which reveal how the system evolves through state-space in time. We found that the awake condition was characterized by both a high-degree of inter-channel interactions, as well as a more flexible, less predictable structure, in contrast to propofol which had less inter-channel interaction, and more predictable, constrained dynamics. Ketamine anesthesia sat between the two extremes, combining elements of both. By combining multiple measures into a sort of “dynamical morphospace”, we can better understand how distinct states relate to each-other.

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] Akeju, O., Westover, M. B., Pavone, K. J., Sampson, A. L., Hartnack, K. E., Brown, E. N., and Purdon, P. L. (2014). Effects of Sevoflurane and Propofol on Frontal Electroencephalogram Power and Coherence. *Anesthesiology*, 121(5):990–998.
- [2] Alkire, M. T., Hudetz, A. G., and Tononi, G. (2008). Consciousness and Anesthesia. *Science (New York, N.Y.)*, 322(5903):876–880.
- [3] Avena-Koenigsberger, A., Goñi, J., Solé, R., and Sporns, O. (2015). Network morphospace. *Journal of The Royal Society Interface*, 12(103):20140881. Publisher: Royal Society.
- [4] Bandt, C. and Pompe, B. (2002). Permutation entropy: a natural complexity measure for time series. *Physical Review Letters*, 88(17):174102.
- [5] 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.
- [6] Carhart-Harris, R. L. (2018). The entropic brain - revisited. *Neuropharmacology*, 142:167–178.
- [7] 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.
- [8] 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.
- [9] Carlsson, G. (2009). Topology and data. *Bulletin of the American Mathematical Society*, 46(2):255–308.
- [10] 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.
- [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] Faber, S. P., Timme, N. M., Beggs, J. M., and Newman, E. L. (2018). Computation is concentrated in rich clubs of local cortical networks. *Network Neuroscience*, pages 1–21.
- [18] Feldman, D. P. and Crutchfield, J. P. (1998). Measures of statistical complexity: Why? *Physics Letters A*, 238(4):244–252.
- [19] Golkowski, D., Larroque, S. K., Vanhaudenhuyse, A., Plenevaux, A., Boly, M., Di Perri, C., Ranft, A., Schneider, G., Laureys, S., Jordan, D., Bonhomme, V., and Ilg, R. (2019). Changes in Whole Brain Dynamics and Connectivity Patterns during Sevoflurane- and Propofol-induced Unconsciousness Identified by Functional Magnetic Resonance Imaging. *Anesthesiology*, 130(6):898–911. Publisher: American Society of Anesthesiologists.
- [20] 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.
- [21] 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.
- [22] 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.
- [23] Hari, R. and Puce, A. (2017). *MEG-EEG Primer*. Oxford University Press.
- [24] 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.
- [25] Hoel, E. and Hoel, E. P. (2017). When the Map Is Better Than the Territory. *Entropy*, 19(5):188.
- [26] 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.

[27] Homayoun, H. and Moghaddam, B. (2007). NMDA Receptor Hypofunction Produces Opposite Effects on Prefrontal Cortex Interneurons and Pyramidal Neurons. *Journal of Neuroscience*, 27(43):11496–11500. Publisher: Society for Neuroscience Section: Brief Communications.

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

[29] Ishizawa, Y., Ahmed, O. J., Patel, S. R., Gale, J. T., Sierra-Mercado, D., Brown, E. N., and Eskandar, E. N. (2016). Dynamics of Propofol-Induced Loss of Consciousness Across Primate Neocortex. *The Journal of Neuroscience*, 36(29):7718–7726.

[30] 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.

[31] 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.

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

[33] 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.

[34] 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.

[35] Lee, U., Kim, S., Noh, G.-J., Choi, B.-M., Hwang, E., and Mashour, G. A. (2009). The directionality and functional organization of frontoparietal connectivity during consciousness and anesthesia in humans. *Consciousness and Cognition: An International Journal*, 18(4):1069–1078. Place: Netherlands Publisher: Elsevier Science.

[36] 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.

[37] 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.

[38] Maksimow, A., Särkelä, M., Långsjö, J. W., Salmi, E., Kaisti, K. K., Yli-Hankala, A., Hinkka-Yli-Salomäki, S., Scheinin, H., and Jääskeläinen, S. K. (2006). Increase in high frequency EEG

activity explains the poor performance of EEG spectral entropy monitor during S-ketamine anesthesia. *Clinical Neurophysiology: Official Journal of the International Federation of Clinical Neurophysiology*, 117(8):1660–1668.

- [39] 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.
- [40] 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.
- [41] McInnes, L., Healy, J., and Melville, J. (2018). UMAP: Uniform Manifold Approximation and Projection for Dimension Reduction. *arXiv:1802.03426 [cs, stat]*. arXiv: 1802.03426.
- [42] Millman, K. J. and Aivazis, M. (2011). Python for Scientists and Engineers. *Computing in Science Engineering*, 13(2):9–12.
- [43] Myers, A., Munch, E., and Khasawneh, F. A. (2019). Persistent homology of complex networks for dynamic state detection. *Physical Review E*, 100(2):022314. Publisher: American Physical Society.
- [44] 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.
- [45] Oliphant, T. E. (2007). Python for Scientific Computing. *Computing in Science Engineering*, 9(3):10–20.
- [46] Patania, A., Petri, G., and Vaccarino, F. (2017). The shape of collaborations. *EPJ Data Science*, 6(1):1–16. Number: 1 Publisher: SpringerOpen.
- [47] 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.
- [48] Perea, J. A. (2018). Topological Time Series Analysis. *arXiv:1812.05143 [cs, math]*. arXiv: 1812.05143.
- [49] 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.
- [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] 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.

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

[53] 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.

[54] 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.

[55] Schroeder, K. E., Irwin, Z. T., Gaidica, M., Nicole Bentley, J., Patil, P. G., Mashour, G. A., and Chestek, C. A. (2016). Disruption of corticocortical information transfer during ketamine anesthesia in the primate brain. *NeuroImage*, 134:459–465.

[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] Timme, N., Alford, W., Flecker, B., and Beggs, J. M. (2014). Synergy, redundancy, and multivariate information measures: an experimentalist’s perspective. *Journal of Computational Neuroscience*, 36(2):119–140.

[59] Tononi, G. (2008). Consciousness as Integrated Information: a Provisional Manifesto. *The Biological Bulletin*, 215(3):216–242.

[60] 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.

[61] 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.

[62] Varley, T. F. (2020). Causal Emergence in Discrete and Continuous Dynamical Systems. *arXiv:2003.13075 [nlin]*. arXiv: 2003.13075.

[63] 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.

[64] 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.

- [65] 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.
- [66] 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).
- [67] 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.
- [68] 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.
- [69] 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.

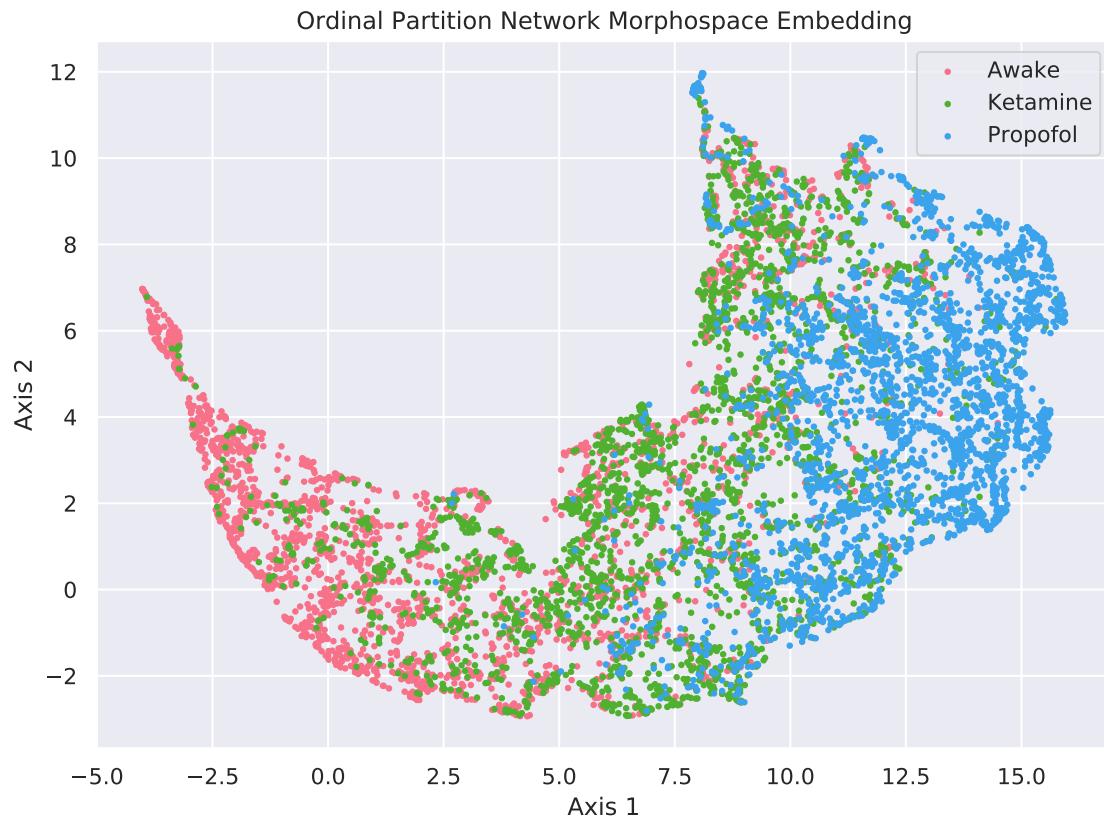


Figure 7: The UMAP [41] embedding of the results of the 5 OPN-based measures for every channel, in each condition. Notice that the ketamine condition forms a kind of boundary between the awake and the propofol conditions, with penetration into both regions. This suggests that, when multiple metrics are taken into consideration, that the state induced by ketamine combines elements of both normal waking consciousness and propofol anaesthesia. Note that, in a UMAP embedding, the resulting axes are not intrinsically meaningful (in contrast to linear embedding algorithms like PCA): the meaningful information is in the distances between individual points, as opposed to the specific values of the coordinates.