

Muting, not fragmentation, of functional brain networks under general anesthesia

Corson N. Areshenkov,^{1,2,*} Joseph Y. Nashed,¹ R. Matthew Hutchison,³ Melina Hutchison,⁴ Ron Levy,^{1,5} Douglas J. Cook,^{1,5} Ravi S. Menon,⁶ Stefan Everling,^{6,7} and Jason P. Gallivan^{1,2}

¹*Centre for Neuroscience Studies, Queens University, Kingston ON, Canada*

²*Department of Psychology, Queens University, Kingston ON, Canada*

³*Biogen, Cambridge MA, USA*

⁴*Massachusetts Eye and Ear Infirmary, Boston MA, USA*

⁵*Department of Surgery, Queens University, Kingston ON, Canada*

⁶*Robarts Research Institute, University of Western Ontario, London ON, Canada*

⁷*Department of Physiology and Pharmacology,*

University of Western Ontario, London ON, Canada

Abstract

Changes in resting-state functional connectivity (rs-FC) under general anesthesia have been widely studied with the goal of identifying neural signatures of consciousness. This work has commonly revealed an apparent fragmentation of whole-brain network structure during unconsciousness, which has been interpreted as reflecting a break-down in connectivity and disruption in the brains ability to integrate information. Here we show, by studying rs-FC under varying depths of isoflurane-induced anesthesia in nonhuman primates, that this apparent fragmentation, rather than reflecting an actual change in network structure, can be simply explained as the result of a global reduction in FC. Specifically, by comparing the actual FC data to surrogate data sets that we derived to test competing hypotheses of how FC changes as a function of dose, we found that increases in whole-brain modularity and the number of network communities considered hallmarks of fragmentation are artifacts of constructing FC networks by thresholding based on correlation magnitude. Taken together, our findings suggest that deepening levels of unconsciousness are instead associated with the increasingly muted expression of functional networks, an observation that constrains current interpretations as to how anesthesia-induced FC changes map onto existing neurobiological theories of consciousness.

* Correspondence email address: areshenk@protonmail.com

I. INTRODUCTION

While much work has focused on how different anesthetics affect ion channels and receptor function at the cellular level (Anis et al. 1983, Franks 2006, Peduto et al. 1991), it remains poorly understood, by comparison, how anesthetics affect the coordinated activity of distributed whole-brain networks (Alkire et al. 2008, Brown et al. 2011). In recent years, resting-state functional MRI (rs-fMRI) has provided important glimpses into the large-scale, network-level effects of anesthesia. This approach, which measures covariance structure in spontaneous low-frequency oscillations in neural activity (Biswal et al. 1995), has repeatedly revealed an apparent fragmentation of functional brain network structure during various states of unconsciousness (Boly et al. 2012b, Hudetz and Mashour 2016, Hutchison et al. 2014). These findings are broadly consistent with work using electroencephalography and electrocorticography reporting a similar breakdown in long-range cortical communication under anesthesia (Lee et al. 2009) and during sleep (Ferrarelli et al. 2010).

Several authors (e.g. Boly et al. 2012b, Hudetz and Mashour 2016, Hutchison et al. 2014, Standage et al. 2019) have, implicitly or explicitly, interpreted these findings through the lens of various neurobiological theories of consciousness, which posit that conscious experience arises through the distributed processing of information throughout the neocortex. For example, the global neuronal workspace theory (Mashour et al. 2020) submits that information is made consciously accessible when it is broadcast widely throughout the cortex by a set of diffusely connected control regions in the prefrontal and parietal cortices. These regions in particular are frequently assigned a key role in the neural substrate of consciousness, and imaging research has revealed a reduction in frontal-parietal connectivity both during sleep (Spoormaker et al. 2012, Tagliazucchi et al. 2013) and under anesthesia (Boly et al. 2012a, Ku et al. 2011). By these, and other similar accounts (e.g. Tononi 2004), the fragmentation of network structure observed during sleep or anesthesia can be viewed as a causal signature of unconsciousness, resulting in a disruption of the brain's ability to integrate information, or to broadcast it widely enough to create conscious awareness (Mashour 2013, Mashour and Hudetz 2018).

Network fragmentation is often concluded on the basis of changes in graph properties (e.g., increases in modularity or the number of communities) estimated from thresholded correlation matrices using a fixed, magnitude threshold. However, in this regime, global decreases – or "muting" – of functional connectivity may also produce the appearance of fragmentation by producing sparser, more disconnected networks (see Figure 1a). As an illustrative example, we contrast work which observed an increase in the modularity and number of communities of significance thresh-

olded networks during sleep (Boly et al. 2012b), with work which found no such effect when using a relative threshold, chosen to obtain a fixed edge density (Uehara et al. 2014). It is difficult to reconcile these results, or to establish that they represent truly distinct phenomena, using graph summary measures alone, as doing so requires a detailed analysis of the underlying network structure. This consideration is particularly pertinent in the case of anesthesia, which has been observed to result in overall decreases in correlation magnitude (Bettinardi et al. 2015, Lv et al. 2016). Further, while several authors have highlighted diminished long-range connectivity under anesthesia – notably between the frontal and parietal cortices (Hudetz and Mashour 2016) – other work has observed similar decreases in local connectivity (Monti et al. 2013). This suggests that anesthetic compounds may affect both short and long range cortical connections, and thus do not simply result in the disconnection of distant cortical regions, but disrupt coordinated neural activity at the local level (Hudetz et al. 2016, Vizuete et al. 2014).

For studies seeking to identify the neural correlates of unconsciousness using functional connectivity, these facts suggest an important distinction – between alterations in network structure on one hand, versus overall changes in correlation magnitude on the other. These different effects can be difficult to disentangle, as evidenced by findings of both (1) an increase in the modularity and number of communities during sleep, and (2) that, despite these network changes, overall network structure remained relatively preserved (Boly et al. 2012b). Given this general ambiguity, as well as the ambiguity as to how such results may support theories of consciousness (Tononi 2004), the goals of the present study were two-fold: First, to ascertain to what extent network fragmentation under anesthesia-induced unconsciousness is attributable to an overall decrease in connectivity strength; and second, to characterize the structure of whole-brain functional connectivity across depths of unconsciousness. To unpack these relationships, we examined changes in rs-fMRI brain network structure in non-human primates across six increasing levels of anesthesia. This allowed us to test for fine-graded changes in network strength and structure across sedation levels, while also assessing critical components of neurobiological theories related to consciousness.

II. RESULTS

A. Reduction in correlation magnitude explains apparent network fragmentation

Figure 2a displays summary statistics for the correlation matrices estimated from each scan. Consistent with previous work (Xie et al. 2019), increasing dose was associated with an overall de-

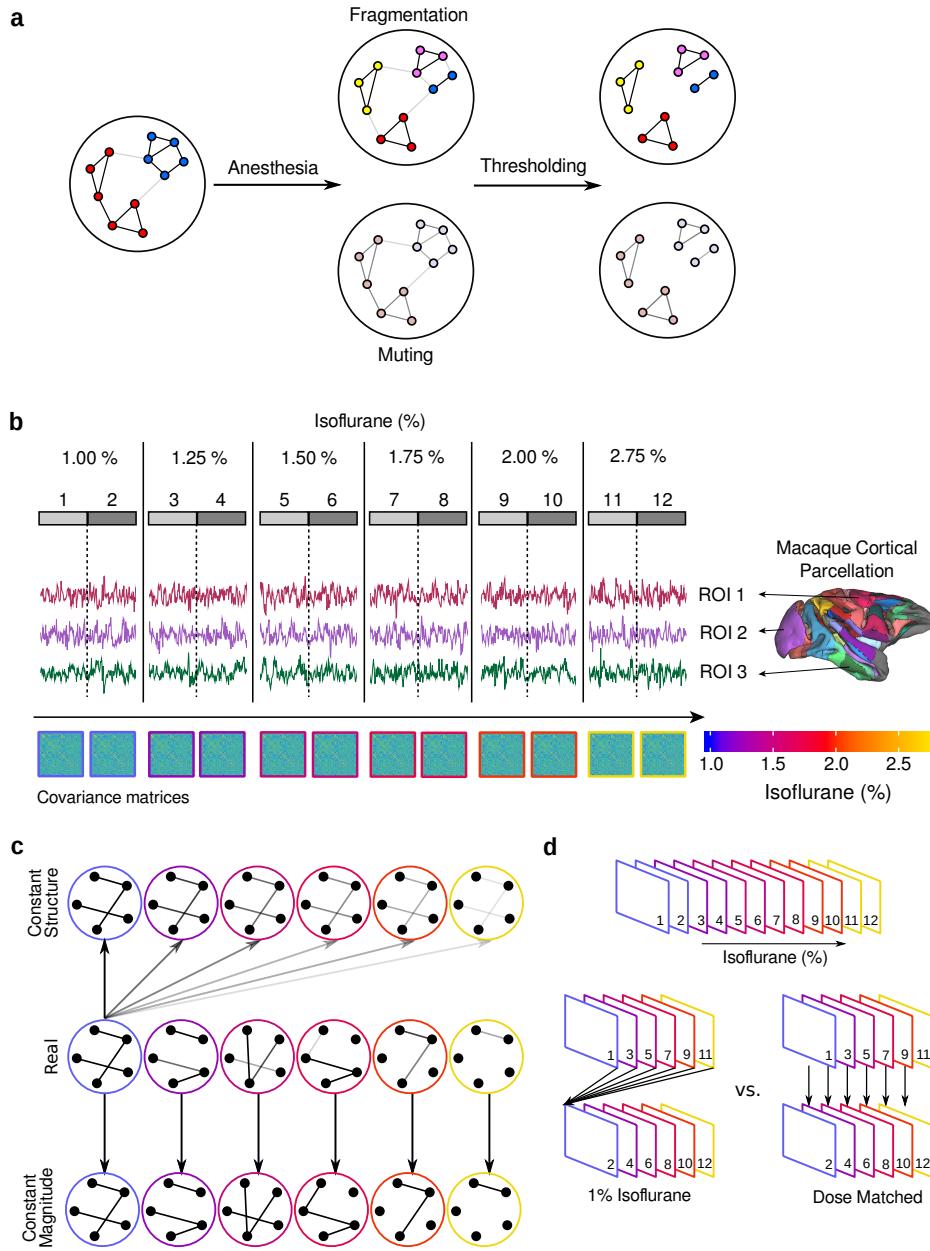


FIG. 1: Competing hypotheses, methods and analytic approaches – **a)** Alternative accounts of the apparent network fragmentation observed during unconsciousness. The *fragmentation* account posits a splitting of conscious brain networks into smaller subnetworks during unconsciousness. An alternative *muting* account explains the apparent fragmentation by a global reduction in correlation magnitude, which results in sparser, more fragmented networks after applying a magnitude threshold. **b)** We collected twelve five-minute resting state scans – two at each of six concentrations of isoflurane. The cortex was parcellated using the LVE atlas (Lewis and Van Essen 2000), and covariance matrices were estimated for each scan. **c)** We studied dose related changes in network statistics both in the **real** data, and in surrogate datasets in which either the correlation structure or magnitude (the mean absolute value of the pairwise correlations) were held constant. In the former (**Constant Structure**), we simply scaled the correlation matrix associated with the lowest dose (1% isoflurane) to match the observed mean magnitude for each scan. In the second (**Constant Magnitude**), we scaled all correlation matrices to have mean magnitude equal to the lowest dose. **d)** To quantify the degree to which the data supported either a constant vs. dose dependent covariance structure, we split each subjects data in two datasets, each comprising one of the two scans at each dose. We then compared the correlation matrices in one half both to the lowest dose (**1% isoflurane**) or to the corresponding dose (**Dose matched**) in the other.

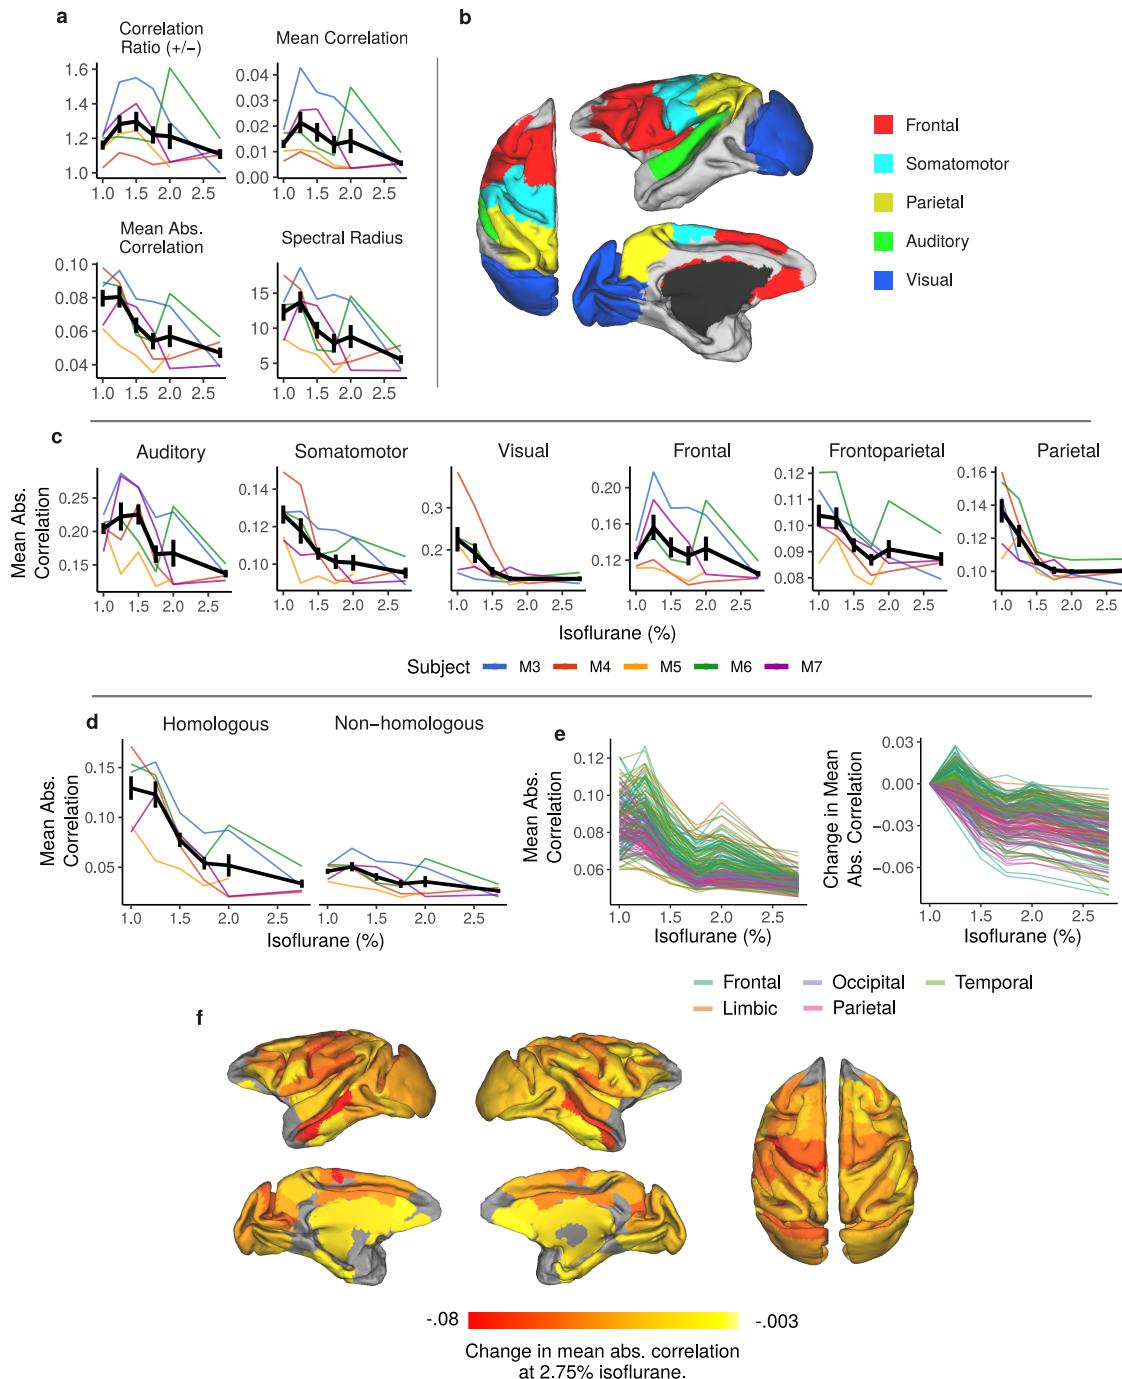


FIG. 2: Correlation magnitude decreases both globally and locally as a function of increasing anesthetic dose – a) Descriptive statistics for correlation matrices. In all figures, colored lines denote means for each subject. Solid black line denotes the mean across all subjects. Error bars are one standard error. **b)** Parcelation of the cortex into primary sensory regions (somatomotor, auditory, and visual), and frontal and parietal cortices. **c)** Mean absolute correlation within each parcel, as well as between the frontal and parietal cortices (frontoparietal). **d)** Mean absolute correlation between bilateral homologues and non-homologous ROIs. **e)** Average muting of functional connectivity per ROI. Figures display the mean absolute correlation, as well as the change in mean absolute correlation relative to the lowest dose (1% isoflurane). **f)** Surface map of the change in mean absolute correlation at 2.75% isoflurane relative to the lowest dose.

crease in correlation magnitude, with no clear change in the ratio of positive to negative correlations

(as in [Bettinardi et al. 2015](#)). We also observed a decrease in the spectral radius (the magnitude of the largest eigenvalue), suggesting an overall loss of low dimensional structure. This reduction in correlation did not appear to be driven by regions in any single network, but was present both in primary sensory and somatomotor regions, as well as in the frontal and parietal cortices, and connections between them (Figures 2b-c). Although bilateral homologs displayed relatively high functional connectivity compared to non-homologous regions at low doses, this connectivity was likewise seen to sharply decrease with increasing dose (Figure 2d). Figures 2e-f displays the decrease in correlation magnitude at the level of individual ROIs, suggesting that all regions show a trend towards decreasing mean connectivity with increasing dose. Together, these effects indicate a global reduction in functional connectivity magnitude as isoflurane dose increases.

Next, to determine the extent to which network changes across dose can be attributed to actual changes in network structure, versus an overall decrease in correlation magnitude, we computed graph summary statistics from the thresholded and binarized correlation matrices, as well as from two surrogate datasets in which either the average correlation structure or the correlation magnitude was held constant across dose. These surrogate datasets were important, as they provided a critical basis for interpreting effects in the real data; if the appearance of network fragmentation is due primarily to an overall decrease in correlation magnitude, then the same pattern of results should be observed when the exact same correlation structure is held constant across doses, and only the magnitude is allowed to vary. Conversely, fragmentation should be abolished when correlation matrices are scaled to have the same average magnitude. We created the **Constant Structure** surrogate dataset by replacing each correlation matrix (from 1.00% - 2.75% isoflurane) with a copy of the average of the subjects two correlation matrices at the lowest dose (1.00%), scaled to have matching correlation magnitude to the real data (defined as the mean absolute value over all pairwise correlations). By contrast, we created the **Constant Magnitude** surrogate dataset by scaling each correlation matrix to have the same average correlation magnitude as in the corresponding subjects lowest dose condition (see Figure 1c). We then thresholded and binarized each correlation matrix using an uncorrected, one-tailed t-test with a significance threshold of $\alpha = .05$. For further comparison, we also binarized the observed correlation matrices using a relative threshold chosen to produce an constant edge density equal to the mean edge density at the lowest dose (**Density Threshold**). If increasing dose is characterized primarily by a overall decrease in correlation strengths, rather than a change in network structure, then graph properties of density thresholded networks should be relatively preserved across dose.

Broadly consistent with previous work ([Boly et al. 2012b](#), [Hutchison et al. 2014](#)) – and the in-

terpretation that brain networks become increasingly fragmented and/or disconnected at increased levels of sedation – we observed an increase in sparsity, the number of communities, modularity, and small-worldness with increasing dose, along with a decrease in network efficiency (Figure 3a). Notably, however, we observed a nearly identical pattern of results in our surrogate dataset in which the correlation structure was held constant across dose (Constant Structure), and only the correlation magnitude was varied. Furthermore, these effects were completely abolished in our surrogate dataset in which the correlations were scaled to have common magnitude across dose, or when the networks were thresholded to achieve a fixed edge density. These results, when taken together, suggest that the observed effect of increasing dose on network statistics (i.e., fragmentation) is better explained as a reduction in overall correlation magnitude, or the muting of constant network structure.

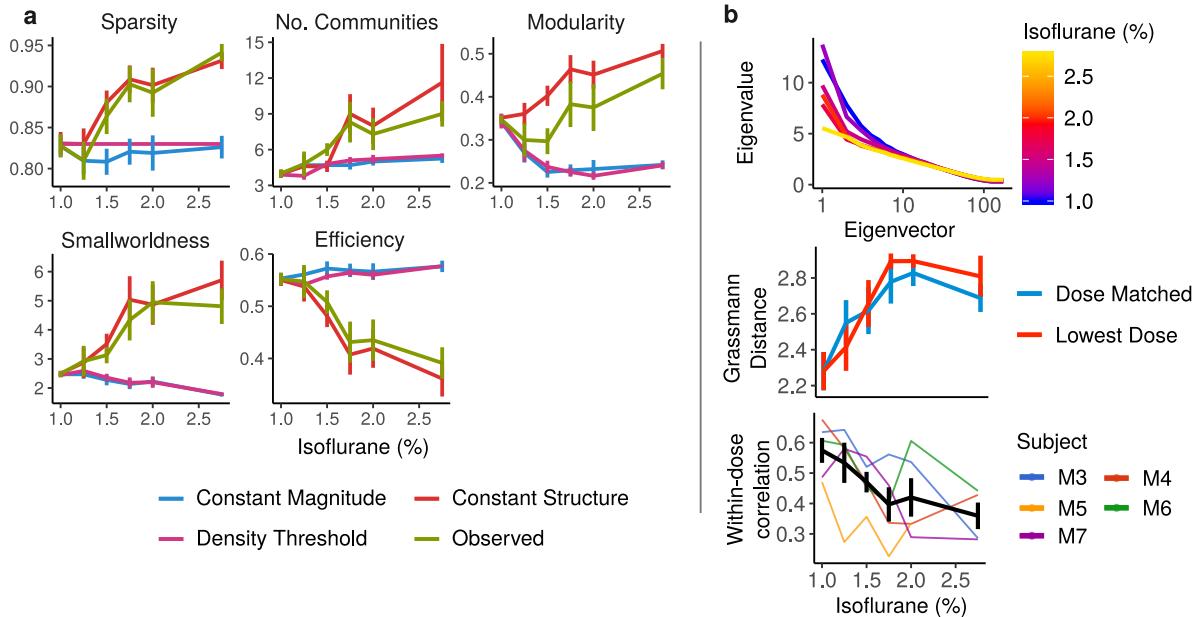


FIG. 3: The appearance of network fragmentation is an artifact of an increasingly muted network structure – a) Network summary statistics. Values denote means across all subjects, while error bars denote one standard error. **Observed** denotes the actual, observed correlation matrices. For **Constant Magnitude**, correlations at each dose level were scaled to match the average magnitude of the correlations at the lowest dose (1% isoflurane). For **Constant Structure**, the correlation matrix for the lowest dose was replicated across all scans, and scaled to match the observed average magnitude. Networks were constructed by thresholding and binarizing correlation matrices by significance using a one-tailed t-test with $\alpha = .05$. For comparison, we also thresholded the observed correlation matrices using a relative threshold to achieve a fixed edge density of .3 (**Density threshold**). Communities were estimated using the Louvain clustering algorithm. Note that the effect of isoflurane on the observed networks is consistent with the scaling of a constant correlation structure. **b)** Mean spectrum of the correlation matrices at each dose (top). The x-axis is shown on a log scale to better display the effects of dose on the leading eigenvalues. After splitting each subjects data into two halves – comprising the first and second scans at each dose, respectively – we compared the correlation matrices in first half either to the corresponding dose (**Dose Matched**), or to the lowest dose (**1% Isoflurane**) in the other. Correlation matrices were compared using the distance between the subspaces spanned by the leading five eigenvectors (middle). The bottom figure displays the correlation between the vectorized correlation matrices from the two scans at each dose, suggesting that the reliability of functional connectivity estimates decreases with increases depth of anesthesia.

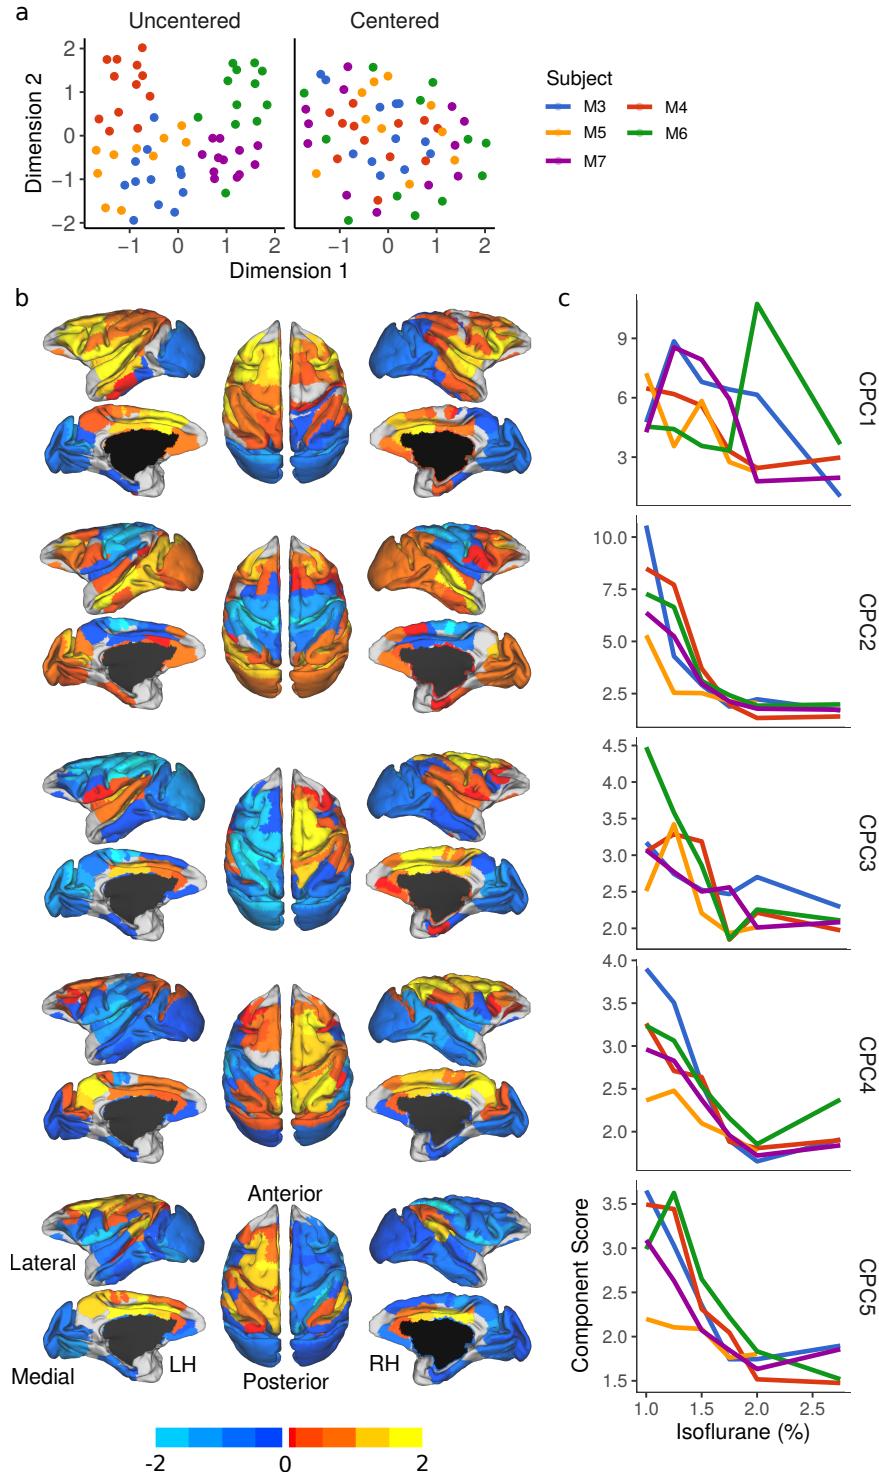


FIG. 4: Latent structure of brain networks is present across all dose levels, but becomes increasingly muted – Common principal component analysis (CPCA) of subjects' centered covariance matrices. The number of displayed components was selected on the examination of the spectrum of the observed correlation matrices (see Figure 3b). Prior to CPCA, subject covariance matrices were centered to remove static subject differences in functional connectivity. **a)** Two-dimensional embedding of subject covariance matrices by uniform manifold approximation (UMAP; McInnes et al. 2018) using a distance matrix constructed by the pairwise geodesic distances. Note the subject level clustering in the uncentered data. **b)** Spatial maps for the top five components. **c)** Component scores for each subject and each scan.

B. Dose effects are well explained by a constant network structure

If the observed dose effects can be explained by the muting of a constant network structure, then the correlation structure at all dose levels should be well approximated by the structure at the lowest dose (1% isoflurane). We directly tested this by splitting each subject's data into two halves, comprising the first and second scans at each dose, respectively. We then compared each whole-brain correlation matrix in the first set either to the corresponding dose in the second, or to the lowest dose, using the distance between the subspaces spanned by the leading eigenvectors (the geodesic distance on the Grassmann manifold; [Edelman et al. 1998](#)). We chose the subspace spanned by the leading five eigenvectors, as this was the range of the spectrum most strongly affected by dose (Figure 3b, top).

Critically, we found a near identical pattern of results in the lowest dose and in the dose matched comparisons (Figure 3b, middle), indicating that the dominant patterns of network structure at higher dosages are well approximated by the structure already present at the lowest dose.

To visualize this structure, we used common principal component analysis (CPCA; [Flury 1984](#), [Trendafilov 2010](#)) to derive a set of components summarizing the correlation structure across dosages. Prior to applying CPCA, subject correlation matrices were centered to remove subject differences in functional connectivity. This decision was motivated by previous findings ([Gratton et al. 2018](#), [Xu et al. 2019](#)) that variability in functional connectivity in Humans and non-human primates is dominated by stable, subject level effects, which may mask the relatively small differences induced by task manipulations. Consistent with these findings, we found strong clustering of the correlation matrices at the subject level (Figure 4a, left) in the uncentered data, which we were able to remove through our centring approach (Figure 4a, right).

Figure 4b-c shows the spatial maps of the components extracted by applying CPCA to the centered correlation matrices, as well as the contributions of these components at each dose. The first of these (CPC1) constitutes a gradient separating visual and medial ventrot temporal areas from parietal, frontal and superior temporal regions. CPC2 constitutes a gradient separating somatomotor areas from the rest of the cortex. CPC3 constitutes a gradient separating right parieto-frontal cortex and bilateral superior temporal cortex from the rest of cortex; CPC5 is a near-mirror image of this same gradient; finally CPC4 constitutes a gradient separating bilateral frontoparietal cortex from the rest of the cortex. The top components are consistent with some of the gradients reported by others in both non-human primates ([Margulies et al. 2016](#), [Yacoub et al. 2020](#)), and in humans ([Hong et al. 2020](#)). Together, these findings support the notion that

a constant network structure is present across all dose levels, and that it is only the expression of this constant network structure that changes across dose. With respect to the latter, indeed we find that these components decreased almost monotonically with increasing dose, asymptoting at approximately 2% isoflurane.

III. DISCUSSION

We analyzed static resting state functional connectivity (rs-FC) under increasing depths of anesthesia in order to characterize dose-related changes in cortical whole-brain network structure. Increasing dose was associated with an increase in network modularity and the number of communities, as well as a decrease in network efficiency, all of which is consistent with the apparent fragmentation discussed in previous literature (Hudetz and Mashour 2016). However, comparisons with surrogate datasets in which either the correlation structure or magnitude were held constant revealed that these effects, rather than reflecting a qualitative change in network structure (i.e., network fragmentation), could be fully explained as an overall reduction in correlation magnitude, which we refer to as muting. Further supporting the idea that network structure was unchanged across dose levels, we showed that the principal components of the functional connectivity matrices at each increasing dose level (from 1.00% - 2.75%) were just as similar to the components at the lowest dose (1.00% isoflurane) as they were to the components in another, dose-matched scan. This suggests that there is no explanatory benefit to assuming a change in correlation structure across dose levels, as this structure can be just as well approximated assuming a fixed structure, already present at the lowest dose. Finally, we used common principal component analysis to derive a set of components that summarize the correlation structure across dose, and show that the expression of these components decreases in a near monotonic fashion as dose increases. Taken together, these findings suggest that deepening levels of unconsciousness are associated with the increasingly muted expression of a constant functional network structure, rather than a break-down or fragmentation of this structure. We discuss both the methodological and theoretical implications of these findings below.

Several neurobiological theories of consciousness center on the brains capacity for information exchange across multiple distributed regions, with unconsciousness being a result of the disruption of this information transmission (Mashour et al. 2020). For instance, the global neuronal workspace model posits that information, initially encoded by a specialized processing unit, is made consciously accessible when it is broadcast widely throughout the cortex, to the many other

specialized processors in the brain (Mashour et al. 2020). By association, conditions that impact this global broadcasting of information, from sleep and anesthesia to brain damage, can result in varying levels of unconsciousness (Mashour and Hudetz 2018). Graph theory measures, such as modularity and network efficiency, have routinely been utilized as indirect estimates of the capacity for information transfer across brain networks, and have often been used to draw inferences concerning changes in network structure during unconsciousness. For instance, a reduction in network efficiency has been observed in both local and global brain networks during anesthetic- and sedative-induced unconsciousness, and has been interpreted as reflecting a dysfunction in communication across the cortex (Hashmi et al. 2017, Monti et al. 2013).

Consistent with this interpretation, decreases in local and global efficiency have also been observed in patients with consciousness disorders, such as unresponsive wakefulness syndrome and the minimally conscious state, and such decreases have been shown to correlate with reductions in patient awareness (Chennu et al. 2014). Likewise, increases in whole-brain modularity and the number of network communities have been observed during both non-rapid eye movement sleep (Boly et al. 2012b) and isoflurane-induced anesthesia (Hutchison et al. 2014, Standage et al. 2019), which has been interpreted as reflecting a literal fragmentation of brain networks into smaller, more isolated processing units. These interpretations are understandably compelling, as graph measures are thought to quantify key aspects of information transmission in the brain, and many of them, such as modularity and efficiency, neatly map onto existing frameworks and hypotheses concerning theories of consciousness (and disruptions thereof).

In the current study, we have demonstrated that changes in graph metrics; such as whole-brain modularity, network efficiency, path length, and number of communities, can all be simply explained by an overall reduction in functional connectivity, which we refer to as muting. This contrasts markedly with the way that changes in graph metrics have been commonly interpreted, which is that they reflect some qualitative change in underlying brain network structure or organization. Distinguishing between these two phenomena (muting vs. fragmentation) requires a detailed analysis of the change in network structure across levels of anesthesia. As we have shown here, graph summary statistics are highly sensitive both to the threshold techniques used to construct the networks, and to overall changes in correlation strength. This makes it difficult to draw firm conclusions about changes in network structure from graph statistics alone. As demonstrated with our surrogate data, the community structure of a thresholded correlation network can be altered in a manner consistent with fragmentation, even when the overall pattern of correlations is held constant. This is consistent with what we observe in the real data: using CPCA, we demon-

strate that, despite a supposed network fractionation (i.e., increased modularity and number of communities), the exact same network structure is present across all dose levels, with only the expression of this structure being weakened (or 'muted') at higher levels of sedation. We also note although we are wary of drawing conclusions from a single observation that subject M5, which experienced an adverse reaction to the anesthetic at the highest dose, showed a relatively flat response to isoflurane, displaying functional connectivity consistent with higher doses even at the lowest dose (Figure 4). This may suggest that hypersensitivity to isoflurane, or the potential for adverse reactions, are detectable in functional connectivity even at low doses.

One of the main observations to emerge in the anesthesia literature, and which has been frequently used to bolster neurobiological theories of consciousness, is that disruptions in consciousness are associated with a break-down in long-range frontal-parietal connectivity. We observed similar reductions in frontal-parietal connectivity under increasing depths of anesthesia; notably, however, this effect was not unique to long-range connectivity, as we also observed comparable decreases within both the frontal and parietal cortices, as well as the primary sensory (visual, auditory) and somatomotor (primary sensory and motor) areas. Taken together, these findings bolster our interpretation that decreases in long-range functional connectivity do not reflect a literal fragmentation of these cortices into distinct networks, as much as it is a consequence of a brain-wide, global reduction in correlation magnitude. This is consistent with previous fundings that propofol anesthesia produces disruptions in local as well as long-range connectivity (Monti et al. 2013). While this seemingly departs from the view that long-range frontal-parietal networks play a unique role in conscious experience *per se*, it nevertheless comports with the broader view, shared by several theories, that unconsciousness stems from a more global disruption of information processing throughout the brain. It is also consistent with views which assign frontal-parietal regions a privileged position due to their role as hub regions facilitating the broadcasting of information widely across the cortex (Mashour et al. 2020), as this broadcasting may be impaired not only by selective disruption of frontal-parietal regions, but also by more global impairments in connectivity.

Our finding that functional connectivity is decreased even in sensory regions is consistent with work showing that anesthetic compounds interfere with neural synchronization even in these regions. For example, local field potentials in the visual cortex form spontaneous, spatially localized patterns of synchronized activity which grow more variable and entropic under anesthesia (Hudetz et al. 2016), and local firing patterns become decorrelated (Vizuete et al. 2014). These local disruptions in coordinated neural activity may result in noise manifesting as a reduction in functional connectivity. Thus, the results we have observed may reflect not simply a reduction in commu-

nication within and between brain networks, but the fading of an existing network structure into a background of neural noise. This account would also be consistent with our observation that increasing depth of anesthesia was associated with a flattening of the spectrum in the estimated correlation matrices, but with relatively little change in the dominant eigenstructure (Figure 3, top and middle). As well as our observation that the reliability of functional connectivity estimates decreased with increasing dose (Figure 3, bottom).

Our findings should be interpreted in light of a few methodological considerations. First, our study did not collect neuroimaging data from monkeys during the awake state, limiting our discussion and interpretations to the network structure present across deepening levels of anesthesia (and presumably unconsciousness). However, [Bettinardi et al. \(2015\)](#) note a gradual emergence of coordinated brain activity during the transition from deep sedation to wakefulness, suggesting that our protocol nonetheless samples a portion of a continuous trajectory spanning waking and unconsciousness. Second, we induced unconsciousness through isoflurane ([Hutchison et al. 2014](#), [Shmuel and Leopold 2008](#), [Vincent et al. 2007](#)), which is one of only several different anesthetics that could have been used. Isoflurane is a potent vasodilator ([Iida et al. 1998](#)), and has been shown to have effects on cerebral blood flow and blood volume ([Li et al. 2013](#), [Masamoto et al. 2006](#)). As such, it is possible that its hemodynamic properties could obscure potential neural changes that occur at higher dose levels. However, studies combining fMRI with electrical recordings, such as EEG-fMRI ([Barttfeld et al. 2015](#), [He et al. 2008](#), [Liu et al. 2013](#), [Ranft et al. 2016](#), [Vincent et al. 2007](#)), have demonstrated a close coupling between changes in neural activity and corresponding changes in functional connectivity. Moreover, our finding of reduced functional connectivity at both the small and large-scale is consistent with similar observations under propofol ([Monti et al. 2013](#)), and studies examining network dynamics under anesthesia have observed similar effects under several different types of anesthesia, from propofol and sevoflurane to ketamine ([Barttfeld et al. 2015](#), [Hutchison et al. 2014](#), [Uhrig et al. 2018](#)). Thus, we think it unlikely that our findings are entirely explainable by the specific mechanisms of action of the anesthetic used here.

Theories relating consciousness to the information processing capacity of the brain must ultimately generate concrete predictions describing the patterns of brain activity supporting conscious awareness. Functional neuroimaging plays a central role in the testing of these predictions, as these techniques allow for the precise characterization of brain activity across various states of consciousness and unconsciousness. Although fMRI affords the opportunity to study the large-scale, structural features of whole-brain networks, the complexity of this data necessitates careful analysis. As we have shown, graph summary measures lack the resolution to fully characterize

whole-brain network structure, or changes in this structure across depths of unconsciousness. Further, they are highly sensitive to the methods used to construct the networks, making it difficult to determine what precise properties of the brain are being quantified by these measures. Our results suggests that a global muting of functional connectivity is a significant feature of isoflurane induced anesthesia, and that this fact is sufficient to explain previously reported results which are commonly attributed to structural changes in whole-brain networks. They also highlight the need to carefully disentangle putative changes in network structure from effects induced by overall changes in correlation magnitude, or by the method of network construction; something that is likely to be difficult through the use of graph summary measures alone. We have not addressed the causal role of this muting, or whether it may simply mask structural network features underlying consciousness; although recent work (Pal et al. 2020) suggests that anesthesia induced unconsciousness can be dissociated from the effects of anesthesia on the cortex, and so the relationship between consciousness and cortical information processing may be more complex than can be decoded from BOLD signal correlations alone.

IV. METHODS

A. Data collection

We reanalyzed data from five Macaque primates (*M. Fascicularis*; 4 female; Mean age 7.8 yrs) collected as part of the experiment reported in Hutchison et al. (2014). All surgical and experimental procedures were carried out in accordance with the Canadian Council of Animal Care policy on the use of laboratory animals and approved by the Animal Use Subcommittee of the University of Western Ontario Council on Animal Care. As data acquisition is thoroughly described by the original authors, we present a more condensed description here. Note however, that our fMRI preprocessing pipeline contains slight differences.

Prior to image acquisition, subjects were injected intramuscularly with atropine (0.4 mg/kg), ipratropium (0.025 mg/kg), and ketamine hydrochloride (7.5 mg/kg), followed by intravenous administration of 3 ml propofol (10 mg/ml) via the saphenous vein. Subjects were then intubated and switched to 1.5% isoflurane mixed with medical air. Each subject was then placed in a custom-built chair and inserted into the magnet bore, at which time the isoflurane level was lowered to 1.00%. Prior to image localization, shimming, and echo-planar imaging (EPI), at least 30 min was allowed for the isoflurane level and global hemodynamics to stabilize at this concentration. We then

acquired 2 functional EPI scans at each of six increasing isoflurane levels: 1.00, 1.25, 1.50, 1.75, 2.00, and 2.75% (0.78, 0.98, 1.17, 1.37, 1.56, and 2.15 minimum alveolar concentration, respectively). We interleaved a 10 min period between each isoflurane level increase to allow for the concentration to stabilize, during which no fMRI data were collected. Throughout the duration of scanning, the monkeys spontaneously ventilated and we monitored physiological parameters (temperature, oxygen saturation, heart rate, respiration, and end-tidal CO₂) to ensure that values were within normal limits. The acquisitions of two anatomical images occurred during the stabilization periods between isoflurane levels.

The monkeys were scanned on an actively shielded 7-Tesla 68-cm horizontal bore scanner with a DirectDrive console (Agilent, Santa Clara, California) with a Siemens AC84 gradient subsystem (Erlangen, Germany). We used a custom in-house conformal five-channel transceive primate-head Radio Frequency (RF) coil. Each functional run consisted of 150 continuous EPI functional volumes (repetition time [TR] = 2000 ms; echo time [TE] = 16 ms; flip angle = 700; slices = 36; matrix = 96 x 96; Field of view [FOV] = 96 x 96 mm²; acquisition voxel size = 1 x 1 x 1 mm³), acquired with GRAPPA = 2. A high-resolution gradient-echo T2 anatomical image was acquired along the same orientation as the functional images (TR = 1100 ms, TE = 8 ms, matrix = 256 x 256, FOV = 96 x 96 mm², acquisition voxel size = 375 x 375 x 1,000 mm³). We also acquired a T1-weighted anatomical image (TE = 2.5 ms, TR = 2300 ms, FOV = 96 x 96 mm², acquisition voxel size = 750 x 750 x 750 mm³).

B. fMRI preprocessing

Functional image preprocessing was implemented using Nipype (Neuroimaging in Python: Pipelines and Interfaces; <http://nipy.org/nipype>). The functional images underwent de-spiking, motion correction (six-parameter affine transformation), and slice-time correction, before brain extraction, and highpass temporal filtering (0.01 Hz; [Hallquist et al. 2013](#), [Power et al. 2014](#)). Functional data was co-registered to its respective T1 anatomical (six degrees of freedom rigid transformation), and then linear (12 degrees of freedom linear affine transformation) and nonlinear transformed to the F99-template ([Van Essen 2002](#)) and parcellated into 174 regions of interest (ROI) using the LVE atlas ([Lewis and Van Essen 2000](#)). All further analyses were done in R (version 3.6.1; [R Core Team 2019](#)).

Nuisance regression was performed using the six motion parameters, their derivatives, their squares, as well as mean white matter and CSF signals, for a total of 26 regressors. We elected

to use the unwhitened residuals rather than performing prewhitening, as the overall residual autocorrelation was minimal (mean Durbin-Watson statistic in the range [1.6-2] for all subjects and all scans), and did not appear to be related to dose.

C. Functional connectivity estimation and comparison

For each scan, we estimated a covariance matrix from the standardized residuals using the shrinkage estimator proposed by [Ledoit and Wolf \(2004\)](#), which has been shown to perform better than the sample covariance matrix in high-dimensional, small-sample settings.

Networks were constructed from the real and surrogate datasets by thresholding the entries of each correlation matrix using a one-tailed t-test with a threshold of $\alpha = .05$. All graph metrics (Figure 3a) save smallworldness were computed from the binarized correlation matrices using the R package igraph ([Csardi and Nepusz 2006](#)), and community detection was performed using the Louvain clustering algorithm ([Blondel et al. 2008](#)). Smallworldness was computed using the R package brainGraph ([Watson 2019](#)).

Structural similarity between correlation matrices was quantified using the distance between the subspaces spanned by the leading five eigenvectors (Figure 3b). These five eigenvectors constitute a basis for a five-dimensional subspace of the full space of 174 ROIs. The set of all such subspaces forms a manifold – called the Grassmann manifold – on which several natural distance measures can be defined ([Edelman et al. 1998](#)). We define the distance between two such subspaces to be the arc length of the geodesic between them, equivalent to the magnitude of the vector of principle angles between the subspaces. Specifically, let \mathbf{V}_1 and \mathbf{V}_2 be matrices whose columns are the top k eigenvectors of the correlation matrices \mathbf{S}_1 and \mathbf{S}_2 , respectively. Then the principal angles between the subspaces spanned by \mathbf{V}_1 and \mathbf{V}_2 are given by $\theta_i = \cos^{-1} \sigma_i$, where $\sigma = \{\sigma_1, \dots, \sigma_k\}$ are the singular values of $\mathbf{V}_1^\top \mathbf{V}_2$. The distance between \mathbf{V}_1 and \mathbf{V}_2 is then

$$d(\mathbf{V}_1, \mathbf{V}_2) = \|\theta\|_2 \tag{1}$$

$$= \sqrt{\sum_{i=1}^k \theta_i^2} \tag{2}$$

D. Centering

To center each subject’s correlation matrices, we took the approach advocated by [Zhao et al. \(2018\)](#), which leverages the natural geometry of the space of covariance matrices. We have

implemented many of the computations required to replicate the analysis in an R package **spdm** (symmetric positive-definite matrix), which is freely available from a Git repository at <https://gitlab.com/fmriToolkit/spdm>.

The procedure is as follows. For each subject i , we computed a geometric mean covariance matrix \bar{S}_i using the fixed-point algorithm described by Congedo et al. (2017), as well as a grand mean \bar{S}_{gm} over all subjects and all scans. We then projected each covariance matrix S_{ij} onto the tangent space at the corresponding subject mean \bar{S}_i to obtain a tangent vector

$$T_{ij} = \bar{S}_i^{1/2} \log(\bar{S}_i^{-1/2} S_{ij} \bar{S}_i^{-1/2}) \bar{S}_i^{1/2}, \quad (3)$$

where \log denotes the matrix logarithm. We then transported each tangent vector to the grand mean \bar{S}_{gm} using the transport proposed by Zhao et al. (2018), obtaining a centered tangent vector

$$T_{ij}^c = GT_{ij}G^\top \quad (4)$$

where $G = \bar{S}_{gm}^{1/2} \bar{S}_i^{-1/2}$. Finally, we projected each centered tangent vector back onto the space of covariance matrices, to obtain the centered covariance matrix

$$S_{ij}^c = \bar{S}_{gm}^{1/2} \exp(\bar{S}_{gm}^{-1/2} T_{ij}^c \bar{S}_{gm}^{-1/2}) \bar{S}_{gm}^{1/2}, \quad (5)$$

where \exp denotes the matrix exponential.

To gain some intuition for this procedure, note that centering a sample of real numbers can be viewed as the process of computing a sample mean, and then applying a translation which takes the sample mean to zero. The resulting centered values can then be viewed as vectors describing the deviations of the observed values from the mean. In the same way, the tangent vectors computed in Eq. 3 can be viewed as the deviations of each correlation matrix from the corresponding subject mean. The transport in Eq. 4 translates these deviations to the grand mean, thus aligning them to a common baseline. A naive approximation to this procedure would involve, for each subject, simply subtracting the mean covariance matrix from the covariance matrix in each scan. This approach has several drawbacks; most importantly, the fact that the difference between two covariance matrices is not necessarily a covariance matrix, and so does not have an obvious interpretation. The procedure we have described above can be viewed as an adaptation of this approach, which respects the geometric structure of the space of covariance matrices.

To visualize the effect of centering (Figure 4a), we derived a two-dimensional embedding of

the set of covariance matrices before and after centering using uniform manifold approximation (UMAP; [McInnes et al. 2018](#)). This procedure was applied to the distance matrices computed from the pairwise distances between covariance matrices ([Smith 2005](#)), defined as

$$d(S_1, S_2) = \sqrt{\sum_{i=1}^n (10 \log_{10} \lambda_i)^2} \quad (6)$$

where $(\lambda_1, \dots, \lambda_n)$ are the generalized eigenvalues of S_1 and S_2 .

E. Common component analysis

After centering, we sought an interpretable, low-dimensional summary of the observed covariance matrices in order to characterize changes in functional connectivity across depths of anesthesia. For a single covariance matrix, this could be accomplished by principal component analysis (PCA), where the eigenvectors of the covariance matrix are used as a basis for a low dimensional subspace capturing the dominant patterns of variability in the BOLD signal observed during a single scan. As we had observations for multiple scans and multiple subjects, we considered two approaches for simultaneously decomposing the full set of covariance matrices.

The first is common principal component analysis (CPCA; [Flury 1984](#), [Trendafilov 2010](#)). The CPCA model attempts to simultaneously diagonalize multiple covariance matrices, and so (informally) assumes that the covariance matrices have identical factor structure, though they may differ in the degree to which they express those factors. As this is a restrictive assumption, we also considered a second model – the common component analysis (CCA) proposed by [Wang et al. \(2011\)](#) – which relaxes the assumption that the set of covariance matrices may be simultaneously diagonalized at the expense of some interpretability. The CPCA model was fit using the R package `cPCA` ([Ziyatdinov et al. 2014](#)), while the CCA model was fit using custom R code implementing the iterative algorithm proposed by [Wang et al. \(2011\)](#). As both models returned highly similar results, we present only the results of CPCA.

V. REFERENCES

Alkire, M. T., Hudetz, A. G., and Tononi, G. (2008). Consciousness and anesthesia. *Science*, 322(5903):876–880.

Anis, N., Berry, S., Burton, N., and Lodge, D. (1983). The dissociative anaesthetics, ketamine and phencyclidine, selectively reduce excitation of central mammalian neurones by n-methyl-aspartate. *British journal of pharmacology*, 79(2):565–575.

Barttfeld, P., Bekinschtein, T. A., Salles, A., Stamatakis, E. A., Adapa, R., Menon, D. K., and Sigman, M. (2015). Factoring the brain signatures of anesthesia concentration and level of arousal across individuals. *NeuroImage: Clinical*, 9:385–391.

Bettinardi, R. G., Tort-Colet, N., Ruiz-Mejias, M., Sanchez-Vives, M. V., and Deco, G. (2015). Gradual emergence of spontaneous correlated brain activity during fading of general anesthesia in rats: Evidences from fmri and local field potentials. *Neuroimage*, 114:185–198.

Biswal, B., Zerrin Yetkin, F., Haughton, V. M., and Hyde, J. S. (1995). Functional connectivity in the motor cortex of resting human brain using echo-planar mri. *Magnetic resonance in medicine*, 34(4):537–541.

Blondel, V. D., Guillaume, J.-L., Lambiotte, R., and Lefebvre, E. (2008). Fast unfolding of communities in large networks. *Journal of statistical mechanics: theory and experiment*, 2008(10):P10008.

Boly, M., Moran, R., Murphy, M., Boveroux, P., Bruno, M.-A., Noirhomme, Q., Ledoux, D., Bonhomme, V., Brichant, J.-F., Tononi, G., et al. (2012a). Connectivity changes underlying spectral eeg changes during propofol-induced loss of consciousness. *Journal of Neuroscience*, 32(20):7082–7090.

Boly, M., Perlberg, V., Marrelec, G., Schabus, M., Laureys, S., Doyon, J., Pélégrini-Issac, M., Maquet, P., and Benali, H. (2012b). Hierarchical clustering of brain activity during human nonrapid eye movement sleep. *Proceedings of the National Academy of Sciences*, 109(15):5856–5861.

Brown, E. N., Purdon, P. L., and Van Dort, C. J. (2011). General anesthesia and altered states of arousal: a systems neuroscience analysis. *Annual review of neuroscience*, 34:601–628.

Chennu, S., Finoia, P., Kamau, E., Allanson, J., Williams, G. B., Monti, M. M., Noreika, V., Arnatkeviciute, A., Canales-Johnson, A., Olivares, F., et al. (2014). Spectral signatures of reorganised brain networks in disorders of consciousness. *PLoS computational biology*, 10(10).

Congedo, M., Barachant, A., and Koopaei, E. K. (2017). Fixed point algorithms for estimating power means of positive definite matrices. *IEEE Transactions on Signal Processing*, 65(9):2211–2220.

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

Edelman, A., Arias, T. A., and Smith, S. T. (1998). The geometry of algorithms with orthogonality constraints. *SIAM journal on Matrix Analysis and Applications*, 20(2):303–353.

Ferrarelli, F., Massimini, M., Sarasso, S., Casali, A., Riedner, B. A., Angelini, G., Tononi, G., and Pearce, R. A. (2010). Breakdown in cortical effective connectivity during midazolam-induced loss of consciousness. *Proceedings of the National Academy of Sciences*, 107(6):2681–2686.

Flury, B. N. (1984). Common principal components in k groups. *Journal of the American Statistical Association*, 79(388):892–898.

Franks, N. P. (2006). Molecular targets underlying general anaesthesia. *British journal of pharmacology*, 147(S1):S72–S81.

Gratton, C., Laumann, T. O., Nielsen, A. N., Greene, D. J., Gordon, E. M., Gilmore, A. W., Nelson, S. M., Coalson, R. S., Snyder, A. Z., Schlaggar, B. L., et al. (2018). Functional brain networks are dominated by stable group and individual factors, not cognitive or daily variation. *Neuron*, 98(2):439–452.

Hallquist, M. N., Hwang, K., and Luna, B. (2013). The nuisance of nuisance regression: spectral misspecification in a common approach to resting-state fmri preprocessing reintroduces noise and obscures functional connectivity. *Neuroimage*, 82:208–225.

Hashmi, J. A., Loggia, M. L., Khan, S., Gao, L., Kim, J., Napadow, V., Brown, E. N., and Akeju, O. (2017). Dexmedetomidine disrupts the local and global efficiency of large-scale brain networks. *Anesthesiology*, 126(3):419.

He, B. J., Snyder, A. Z., Zempel, J. M., Smyth, M. D., and Raichle, M. E. (2008). Electrophysiological correlates of the brain’s intrinsic large-scale functional architecture. *Proceedings of the National Academy of Sciences*, 105(41):16039–16044.

Hong, S.-J., Xu, T., Nikolaidis, A., Smallwood, J., Margulies, D. S., Bernhardt, B., Vogelstein, J., and Milham, M. (2020). Toward a connectivity gradient-based framework for reproducible biomarker discovery. *bioRxiv*.

Hudetz, A. G. and Mashour, G. A. (2016). Disconnecting consciousness: Is there a common anesthetic end-point? *Anesthesia and analgesia*, 123(5):1228.

Hudetz, A. G., Vizuete, J. A., Pillay, S., and Mashour, G. A. (2016). Repertoire of mesoscopic cortical activity is not reduced during anesthesia. *Neuroscience*, 339:402–417.

Hutchison, R. M., Hutchison, M., Manning, K. Y., Menon, R. S., and Everling, S. (2014). Isoflurane

induces dose-dependent alterations in the cortical connectivity profiles and dynamic properties of the brain's functional architecture. *Human brain mapping*, 35(12):5754–5775.

Iida, H., Ohata, H., Iida, M., Watanabe, Y., and Dohi, S. (1998). Isoflurane and sevoflurane induce vasodilation of cerebral vessels via atp-sensitive k⁺ channel activation. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 89(4):954–960.

Ku, S.-W., Lee, U., Noh, G.-J., Jun, I.-G., and Mashour, G. A. (2011). Preferential inhibition of frontal-to-parietal feedback connectivity is a neurophysiologic correlate of general anesthesia in surgical patients. *PLoS one*, 6(10).

Ledoit, O. and Wolf, M. (2004). A well-conditioned estimator for large-dimensional covariance matrices. *Journal of multivariate analysis*, 88(2):365–411.

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*, 18(4):1069–1078.

Lewis, J. W. and Van Essen, D. C. (2000). Corticocortical connections of visual, sensorimotor, and multimodal processing areas in the parietal lobe of the macaque monkey. *Journal of Comparative Neurology*, 428(1):112–137.

Li, C.-X., Patel, S., Auerbach, E. J., and Zhang, X. (2013). Dose-dependent effect of isoflurane on regional cerebral blood flow in anesthetized macaque monkeys. *Neuroscience letters*, 541:58–62.

Liu, X., Pillay, S., Li, R., Vizuete, J. A., Pechman, K. R., Schmainda, K. M., and Hudetz, A. G. (2013). Multiphasic modification of intrinsic functional connectivity of the rat brain during increasing levels of propofol. *Neuroimage*, 83:581–592.

Lv, P., Xiao, Y., Liu, B., Wang, Y., Zhang, X., Sun, H., Li, F., Yao, L., Zhang, W., Liu, L., et al. (2016). Dose-dependent effects of isoflurane on regional activity and neural network function: a resting-state fmri study of 14 rhesus monkeys: an observational study. *Neuroscience letters*, 611:116–122.

Margulies, D. S., Ghosh, S. S., Goulas, A., Falkiewicz, M., Huntenburg, J. M., Langs, G., Bezgin, G., Eickhoff, S. B., Castellanos, F. X., Petrides, M., et al. (2016). Situating the default-mode network along a principal gradient of macroscale cortical organization. *Proceedings of the National Academy of Sciences*, 113(44):12574–12579.

Masamoto, K., Kim, T., Fukuda, M., Wang, P., and Kim, S.-G. (2006). Relationship between neural, vascular, and bold signals in isoflurane-anesthetized rat somatosensory cortex. *Cerebral cortex*, 17(4):942–950.

Mashour, G. A. (2013). Cognitive unbinding: a neuroscientific paradigm of general anesthesia and related

states of unconsciousness. *Neuroscience & Biobehavioral Reviews*, 37(10):2751–2759.

Mashour, G. A. and Hudetz, A. G. (2018). Neural correlates of unconsciousness in large-scale brain networks. *Trends in neurosciences*, 41(3):150–160.

Mashour, G. A., Roelfsema, P., Changeux, J.-P., and Dehaene, S. (2020). Conscious processing and the global neuronal workspace hypothesis. *Neuron*, 105(5):776–798.

McInnes, L., Healy, J., and Melville, J. (2018). Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*.

Monti, M. M., Lutkenhoff, E. S., Rubinov, M., Boveroux, P., Vanhaudenhuyse, A., Gosseries, O., Bruno, M.-A., Noirhomme, Q., Boly, M., and Laureys, S. (2013). Dynamic change of global and local information processing in propofol-induced loss and recovery of consciousness. *PLoS computational biology*, 9(10).

Pal, D., Li, D., Dean, J. G., Brito, M. A., Liu, T., Fryzel, A. M., Hudetz, A. G., and Mashour, G. A. (2020). Level of consciousness is dissociable from electroencephalographic measures of cortical connectivity, slow oscillations, and complexity. *Journal of Neuroscience*, 40(3):605–618.

Peduto, V., Concas, A., Santoro, G., Biggio, G., and Gessa, G. (1991). Biochemical and electrophysiologic evidence that propofol enhances gabaergic transmission in the rat brain. *Anesthesiology*, 75(6):1000–1009.

Power, J. D., Schlaggar, B. L., and Petersen, S. E. (2014). Studying brain organization via spontaneous fmri signal. *Neuron*, 84(4):681–696.

R Core Team (2019). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.

Ranft, A., Golkowski, D., Kiel, T., Riedl, V., Kohl, P., Rohrer, G., Pientka, J., Berger, S., Thul, A., Maurer, M., et al. (2016). Neural correlates of sevoflurane-induced unconsciousness identified by simultaneous functional magnetic resonance imaging and electroencephalography. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 125(5):861–872.

Shmuel, A. and Leopold, D. A. (2008). Neuronal correlates of spontaneous fluctuations in fmri signals in monkey visual cortex: implications for functional connectivity at rest. *Human brain mapping*, 29(7):751–761.

Smith, S. T. (2005). Covariance, subspace, and intrinsic cramer-rao bounds. *IEEE Transactions on Signal Processing*, 53(5):1610–1630.

Spoormaker, V. I., Gleiser, P., and Czisch, M. (2012). Frontoparietal connectivity and hierarchical structure of the brains functional network during sleep. *Frontiers in neurology*, 3:80.

Standage, D., Areshenkov, C. N., Nashed, J. Y., Hutchison, R. M., Hutchison, M., Heinke, D., Menon, R. S., Everling, S., and Gallivan, J. P. (2019). Dynamic reconfiguration, fragmentation and integration of

whole-brain modular structure across depths of unconsciousness. *bioRxiv*, page 783175.

Tagliazucchi, E., von Wegner, F., Morzelewski, A., Brodbeck, V., Jahnke, K., and Laufs, H. (2013). Breakdown of long-range temporal dependence in default mode and attention networks during deep sleep. *Proceedings of the National Academy of Sciences*, 110(38):15419–15424.

Tononi, G. (2004). An information integration theory of consciousness. *BMC neuroscience*, 5(1):42.

Trendafilov, N. T. (2010). Stepwise estimation of common principal components. *Computational Statistics & Data Analysis*, 54(12):3446–3457.

Uehara, T., Yamasaki, T., Okamoto, T., Koike, T., Kan, S., Miyauchi, S., Kira, J.-i., and Tobimatsu, S. (2014). Efficiency of a small-world brain network depends on consciousness level: a resting-state fmri study. *Cerebral cortex*, 24(6):1529–1539.

Uhrig, L., Sitt, J. D., Jacob, A., Tasserie, J., Barttfeld, P., Dupont, M., Dehaene, S., and Jarraya, B. (2018). Resting-state dynamics as a cortical signature of anesthesia in monkeys. *Anesthesiology: The Journal of the American Society of Anesthesiologists*, 129(5):942–958.

Van Essen, D. C. (2002). Windows on the brain: the emerging role of atlases and databases in neuroscience. *Current opinion in neurobiology*, 12(5):574–579.

Vincent, J. L., Patel, G. H., Fox, M. D., Snyder, A. Z., Baker, J. T., Van Essen, D. C., Zempel, J. M., Snyder, L. H., Corbetta, M., and Raichle, M. E. (2007). Intrinsic functional architecture in the anaesthetized monkey brain. *Nature*, 447(7140):83–86.

Vizuete, J. A., Pillay, S., Ropella, K. M., and Hudetz, A. G. (2014). Graded defragmentation of cortical neuronal firing during recovery of consciousness in rats. *Neuroscience*, 275:340–351.

Wang, H., Banerjee, A., and Boley, D. (2011). Common component analysis for multiple covariance matrices. In *Proceedings of the 17th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 956–964. ACM.

Watson, C. G. (2019). *brainGraph: Graph Theory Analysis of Brain MRI Data*. R package version 2.7.3.

Xie, H., Chung, D. Y., Kura, S., Sugimoto, K., Aykan, S. A., Wu, Y., Sakadžić, S., Yaseen, M. A., Boas, D. A., and Ayata, C. (2019). Differential effects of anesthetics on resting state functional connectivity in the mouse. *Journal of Cerebral Blood Flow & Metabolism*, page 0271678X19847123.

Xu, T., Sturgeon, D., Ramirez, J. S., Froudast-Walsh, S., Margulies, D. S., Schroeder, C. E., Fair, D. A., and Milham, M. P. (2019). Interindividual variability of functional connectivity in awake and anesthetized rhesus macaque monkeys. *Biological Psychiatry: Cognitive Neuroscience and Neuroimaging*, 4(6):543–553.

Yacoub, E., Grier, M. D., Auerbach, E. J., Lagore, R. L., Harel, N., Ugurbil, K., Adriany, G., Zilverstand,

A., Hayden, B. Y., Heilbronner, S. R., et al. (2020). Ultra-high field (10.5 t) resting state fmri in the macaque. *bioRxiv*.

Zhao, Q., Kwon, D., and Pohl, K. M. (2018). A riemannian framework for longitudinal analysis of resting-state functional connectivity. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 145–153. Springer.

Ziyatdinov, A., Kanaan-Izquierdo, S., Trendafilov, N. T., and Perera-Lluna, A. (2014). *cPCA: Methods to perform Common Principal Component Analysis (CPCA)*. R package version 0.1.2.