

Recovery of consciousness is mediated by a network of discrete metastable activity states

Andrew E. Hudson^{a,1}, Diany Paola Calderon^{b,1}, Donald W. Pfaff^{b,2}, and Alex Proekt^{b,c,2}

^aDepartment of Anesthesiology and Perioperative Medicine, David Geffen School of Medicine, University of California, Los Angeles, CA 90095; ^bLaboratory for Neurobiology and Behavior, The Rockefeller University, New York, NY 10065; and ^cDepartment of Anesthesiology, Weill Cornell Medical College, New York, NY 10021

Contributed by Donald W. Pfaff, May 13, 2014 (sent for review January 15, 2014)

It is not clear how, after a large perturbation, the brain explores the vast space of potential neuronal activity states to recover those compatible with consciousness. Here, we analyze recovery from pharmacologically induced coma to show that neuronal activity en route to consciousness is confined to a low-dimensional subspace. In this subspace, neuronal activity forms discrete metastable states persistent on the scale of minutes. The network of transitions that links these metastable states is structured such that some states form hubs that connect groups of otherwise disconnected states. Although many paths through the network are possible, to ultimately enter the activity state compatible with consciousness, the brain must first pass through these hubs in an orderly fashion. This organization of metastable states, along with dramatic dimensionality reduction, significantly simplifies the task of sampling the parameter space to recover the state consistent with wakefulness on a physiologically relevant timescale.

anesthesia | state transitions | spectral analysis | emergence

The brain exhibits a remarkable ability to recover normal function associated with wakefulness, even after large perturbations to its activity. Two well-known examples of this are anesthesia and brain injury (1, 2). How the brain recovers from large perturbations currently is unknown. Given the number of neurons involved, the potential space of activity is huge. Thus, it is not clear how the brain samples the vast parameter space to discover patterns of activity that are consistent with consciousness after a large perturbation.

The simplest possibility for the recovery of consciousness (ROC) is that, driven by noise inherent in many aspects of neuronal activity (3), the brain performs a random walk through the parameter space until it eventually enters the region that is consistent with consciousness. An alternative possibility is that although the motion through the parameter space is not random, the trajectory nonetheless is smooth. Lastly, it is possible that en route to ROC, the brain passes through a set of discrete metastable states—that is, a series of jumps between long-lived activity configurations.

The utility of metastable intermediates to the problem of ROC is well illustrated by analogy with protein folding. Levinthal's paradox (4) refers to the implausibility of a denatured protein recovering its native fold conformation by random walk alone, as the time required to randomly explore the conformational space will rapidly exceed the age of the universe, even for a small number of residues. However, energetically favorable metastable intermediate states allow denatured proteins to assume their native conformation rapidly. Thus, we hypothesized that after large perturbations, brain dynamics during ROC are structured into discrete metastable intermediate states.

If metastable intermediate states do exist, transitions between them must be considered. It is unclear a priori, for example, whether there will be an obligate intermediate state that must occur en route to consciousness, or if many different routes through intermediate states enable ROC. In this work, we approximate transitions between metastable intermediate states as

Markovian—dependent only on the current state of the system—so that characterizing the transition probabilities between states sufficiently characterizes the network of metastable intermediate states. Several examples of possible network structures are (i) an ordered “chain” in which each state connects to only two others; (ii) a “small-world” structure, in which most states are connected only locally whereas a few central hub states connect widely, allowing rapid long-distance travel through the network; and (iii) a lattice structure, in which all states have approximately the same connectivity, allowing multiple routes to ROC.

In this report, we demonstrate that in rats under isoflurane anesthesia, ROC occurs after the brain traverses a series of metastable intermediate activity configurations. We demonstrate that the recovery process is not compatible with a random walk or another continuous process, nor does it occur as a single jump. A low-dimensional subspace allows visualization of key features of the recovery process, including clusters of activity consistent with metastable intermediates. These clusters of activity have structured transition properties such that only certain transitions are observed en route to ROC, suggesting that certain states function as hubs.

Results

To analyze the dynamics of ROC, we simultaneously recorded local field potentials (LFPs) from the anterior cingulate and retrosplenial cortices and the intralaminar thalamus (Fig. S1) in rats ($n = 6$) during recovery from general anesthesia induced with isoflurane. These interconnected areas are involved in brain

Significance

How does the brain recover consciousness after significant perturbations such as anesthesia? The simplest answer is that as the anesthetic washes out, the brain follows a steady and monotonic path toward consciousness. We show that this simple intuition is incorrect. We varied the anesthetic concentration to parametrically control the magnitude of perturbation to brain dynamics while analyzing the characteristics of neuronal activity during recovery of consciousness. We find that, en route to consciousness, the brain passes through several discrete activity states. Although transitions between certain of these activity states occur spontaneously, transitions between others are not observed. Thus, the network formed by these state transitions gives rise to an ordered sequence of states that mediates recovery of consciousness.

Author contributions: A.E.H., D.P.C., D.W.P., and A.P. designed research; A.E.H., D.P.C., and A.P. performed research; A.E.H. and A.P. analyzed data; and A.E.H., D.P.C., D.W.P., and A.P. wrote the paper.

The authors declare no conflict of interest.

Freely available online through the PNAS open access option.

¹A.E.H. and D.P.C. contributed equally to this work.

²To whom correspondence may be addressed. E-mail: proekt@gmail.com or pfaff@rockefeller.edu.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1408296111/-DCSupplemental.

arousal and anesthesia (5, 6). The power spectra of the LFPs quantify the distribution of signal power among different frequencies and provide a convenient and statistically robust (7) description of patterns of activity that has been used extensively (e.g., refs. 8, 9) to distinguish neuronal activity in the awake and inactivated brain (e.g., Fig. 1). Thus, in what follows, we chose to quantify brain activity in terms of its spectrum.

We used isoflurane to elicit burst suppression, because its slow pharmacokinetics (10) allowed us to focus on the intrinsic brain dynamics rather than on the kinetics of anesthetic washout. To ensure that all of our experiments began with comparable magnitude perturbation to brain activity, we began each series of experiments with an isoflurane concentration of 1.75%, which reliably produced burst suppression, a pathological pattern of activity seen after trauma (2), anesthesia (11), hypothermia (12), encephalopathy (13), hypoxia (14), and others (e.g., Fig. 1A, blue trace). Burst suppression is defined by episodic low-frequency oscillations (bursts) punctuated by periods of quiescence (suppression) in the electroencephalogram (EEG) and LFPs that correlate with synchronous depolarization of cortical neurons and electrical silence of neuronal membranes (15), respectively. Any further inactivation of the brain results in persistent electrical quiescence. In the awake brain, conversely, persistent high-frequency low-amplitude oscillations (e.g., Fig. 1A, red trace) corresponding to asynchronous neuronal firing (16, 17) are observed.

Animals were maintained at a fixed anesthetic concentration for at least 1 h, after which the concentration was decreased by 0.25% until ROC (usually occurring at 0.75%), defined as the onset of spontaneous movement of the limbs and postural muscles (*SI Materials and Methods*). Although the onset of movement is an imperfect measure, we chose it as an endpoint for several reasons: (i) Onset of limb movement can be detected

readily. (ii) The anesthetic concentration at which humans lose consciousness is correlated closely with the anesthetic concentration at which experimental animals lose their righting reflex (reviewed in ref. 18). (iii) There is no single accepted measure that reliably detects onset of consciousness based on brain activity. (iv) Onset of movement is a conservative estimate of the onset of consciousness in that in the absence of brainstem lesion, it is unlikely that the animal will be awake and not moving during emergence from a pure volatile anesthetic (note that use of an opiate would complicate this, as the animal might be awake but not moving).

The slow titration of isoflurane allowed a prolonged sampling of each anesthetic concentration at steady state. While we controlled inspired anesthetic concentration to make sure that fluctuations in the respiratory dynamics did not result in fluctuations in the brain anesthetic concentration, we monitored respiratory rate (*SI Materials and Methods*). We could not detect statistically significant changes in respiratory rate during fixed anesthetic exposure (repeated measures ANOVA, $df = 19$, $F = 0.672$, $P = 0.830$). Thus, given no change in tidal volume, the brain anesthetic concentration likely will remain constant for a large fraction of the time exposed to a fixed inspired anesthetic concentration.

ROC Is Not Consistent with a Random Walk—Even with Constraints.

Although the characteristics of neuronal activity in the anesthetized and awake brain are well known, how the brain navigates between these states is less clear. Many aspects of neuronal dynamics are stochastic (3). Unsurprisingly, changes in the spectrum from one temporal window to the next are well approximated by multidimensional uncorrelated noise (Fig. S2). This is consistent with the simplest null hypothesis that on a fast time scale (1-s step between consecutive spectral windows), neuronal dynamics perform a random walk. However, even a constrained random walk using the observed pairwise differences between spectra as steps (*SI Materials and Methods*) fails to reliably reach patterns of activity consistent with wakefulness (Fig. 1C). Considering more aspects of neuronal activity exacerbates this problem, as the return of a random walker is guaranteed in only two dimensions at most (19). Thus, to attain ROC on a physiologically relevant time scale, the neuronal activity must be structured.

Indeed, while the anesthetic was decreased slowly and monotonically, neuronal activity switched abruptly between several distinct modes that persisted on the scale of minutes (Fig. 2 spectra; Fig. S3 traces). These fluctuations, evidenced by abrupt changes in power, appear simultaneously in anatomically separated brain regions, signifying a global change in the dynamics of the extended thalamocortical networks. Remarkably, there is no one-to-one correspondence between brain activity and anesthetic concentration—several patterns are seen at a single concentration. These state transitions reveal the essential metastable intermediates produced by the brain en route to ROC.

A Low-Dimensional Subspace Captures Significant Dynamics of ROC.

Correlated fluctuations in spectral power at different anatomical locations suggest that the dynamics of recovery are embedded in a low-dimensional subspace. To analyze this subspace, we first encoded brain activity at time t as point $X(t) = \{x_1, \dots, x_n\}$ in a multidimensional space where each element x_i corresponds to the fraction of power contained at i th frequency concatenated across multiple simultaneously recorded channels during a time window centered at t (*SI Materials and Methods*). We then performed dimensionality reduction of the matrix containing the evolution of brain activity encoded in this fashion using principal component analysis (PCA; *SI Materials and Methods*). PCA exploits the covariance structure of the variables, in this case distribution of power among different frequencies in different anatomical regions, to identify mutually orthogonal directions—principal components (PCs)—formed by linear combinations of

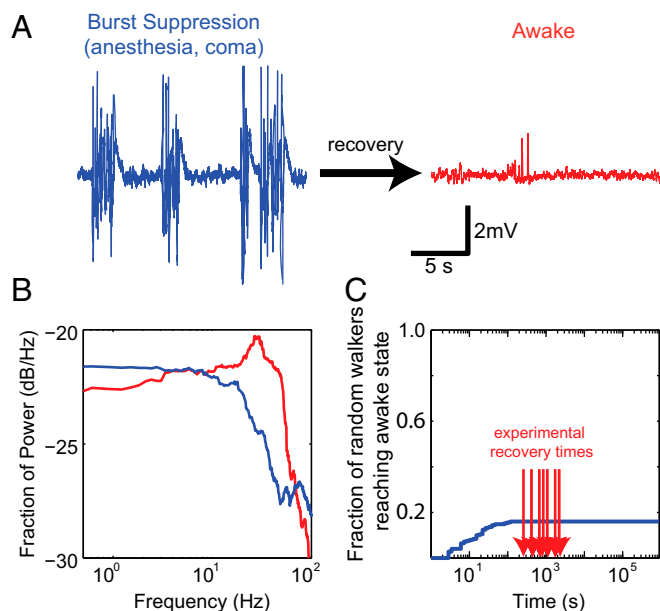


Fig. 1. ROC is not attainable by random walk. (A) Cortical LFP exemplifying burst suppression (blue) observed in pathological states (e.g., coma, anesthesia). LFP observed in the awake brain is shown in red. (B) The power spectra for the traces in A and B (blue and red, respectively) distinguish these activity patterns in the frequency domain. Power contained at each frequency is expressed as the fraction of total power. Differences between the spectra are distributed among many frequencies. (C) Cumulative distribution of recovery times of random walk simulations (*SI Materials and Methods*) shows the improbability of recovery by random walk alone. Red arrows show the experimentally observed recovery times.

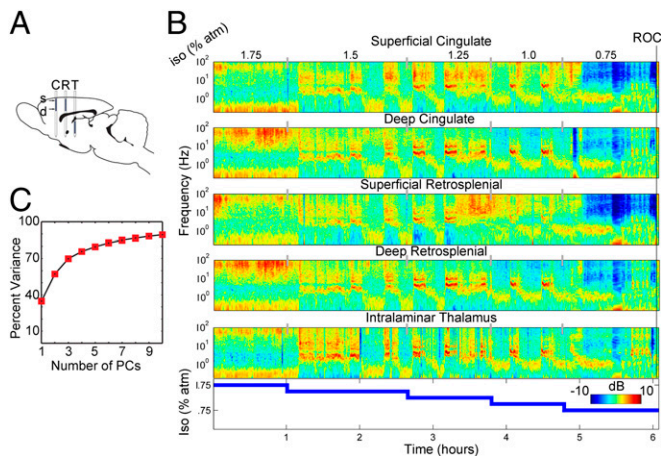


Fig. 2. Time-resolved spectrograms reveal state transitions (A) Diagram of the multi-electrode array used to record simultaneous activity in the anterior cingulate (C) and retrosplenial (R) cortices, as well as the intralaminar thalamus (T), superimposed on the sagittal brain section. (B) Time–frequency spectrograms at different anatomical locations during ROC. The power spectral density at each point in time–frequency space indicates the deviation from the mean spectrum on a decibel color scale as the anesthetic concentration is decreased (Bottom) from 1.75% to 0.75% in 0.25% increments until ROC. (C) Data of the kind shown in B pooled across all animals and all anesthetic concentrations were subjected to PCA (*SI Materials and Methods*). Percent of variance is plotted as a function of the number of PCs. Dynamics of ROC largely are confined to a 3D subspace.

the original variables along which most of the fluctuations occur. Using this approach, we captured $\sim 70\%$ of the variance in just three dimensions (a reduction from 1,245 dimensions; *SI Materials and Methods*) (Fig. 2C). This dimensionality reduction greatly simplifies the recovery from a perturbation.

The position of the data in the 3D subspace spanned by the first three PCs is determined by the similarity of the spectrum to each of the three PCs. For example, the spectrum most similar in shape to PC1 will have the highest coordinate along that

dimension. The shapes of the PCs (Fig. 3A), therefore, indicate the ranges of frequencies in which correlated fluctuations occur in different layers of the cortex and in the thalamus.

Consistent with the laminar architecture of the cortex, PCs demonstrate a laminar pattern (Fig. 3A)—superficial and deep cortical layers form two distinct groups. Although the intralaminar thalamus contains neurons that project to the superficial cortical layers (20), the behavior of the thalamus is distinct from that of superficial cortical layers. For example, the second PC in the thalamus closely resembles the third PC in the superficial cortical layers in that it emphasizes an increase in the power of high-frequency oscillations usually associated with increased arousal. The fact that this increase in high-frequency activity is present in orthogonal PCs implies that activation of the thalamus is separable from activation of the cortex.

Dimensionality reduction (Figs. 2 and 3) was performed on the dataset concatenated across all animals (*Materials and Methods*). To make sure the observed dimensionality reduction was not an artifact of the concatenation, we subjected the data from each animal taken individually to PCA in the same way as for Figs. 2 and 3 (Fig. S4). The dimensionality reduction in each animal is comparable to that in the concatenated dataset. The PCs obtained in each animal and those in the concatenated dataset are not expected to be identical. Furthermore, truncation of the PCA after the first three dimensions is a highly nonlinear operator. Thus, to make sure the concatenation did not introduce dramatic differences in the structure of the data obtained in each experiment, we correlated distances between points in the animal-based and combined PCA (Fig. S4B and C). In all cases, the distances in the animal-based and combined PCAs were highly correlated. Thus, although concatenation may result in the rotation or stretching of the dataset, it does not strongly affect the interrelationship between points obtained in each experiment individually.

Note the key distinction between the results in Figs. 2C and 3 and those in Fig. S2. To characterize the dynamics of recovery from anesthesia, both position—i.e., activity—and velocity—i.e., change in activity—must be considered. Whereas in Figs. 2C and 3

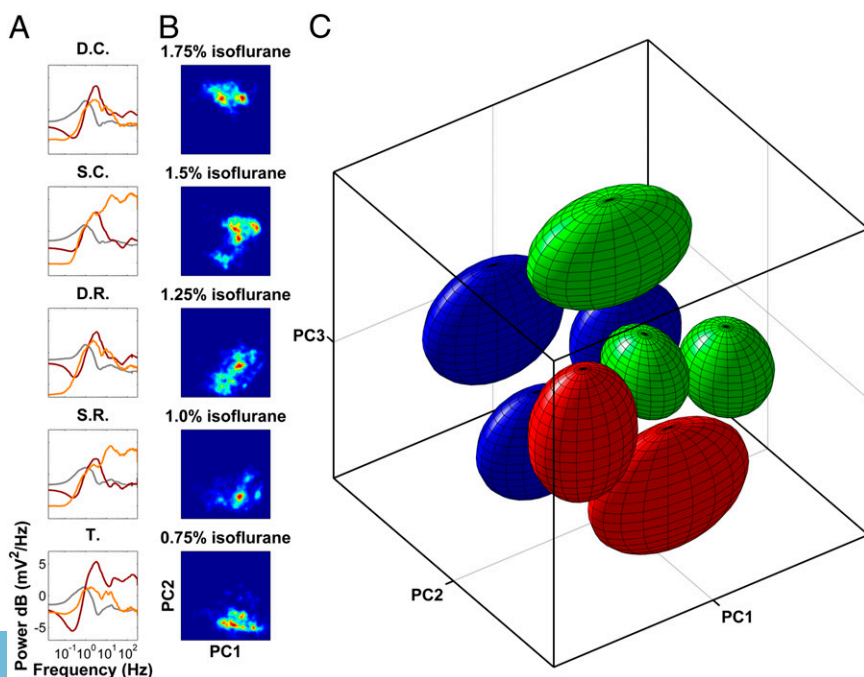


Fig. 3. ROC is characterized by individually stabilized, discrete activity patterns. (A) PC1, -2, and -3 (gray, burgundy, and orange) plotted as a function of frequency and projected onto the corresponding anatomical sites. PCs reveal laminar cortical architecture whereby superficial and deep cortical layers form two distinct groups. High-frequency oscillations are captured by PC2 in the thalamus and PC3 in the superficial cortical layers. Thus, activation of neuronal activity in the thalamus is separable from that in the cortex. D.C., deep cingulate; D.R., deep retrosplenial; S.C., superficial cingulate; S.R., superficial retrosplenial; T, thalamus. (B) Probability density of data from all animals projected onto the plane spanned by PC1 and PC2 (red shows increased probability) shows multiple distinct peaks that change in prevalence and location, depending on anesthetic concentration. (C) In the space spanned by the first three PCs, data form eight distinct clusters (*SI Materials and Methods*). The approximate location of each cluster is shown by an ellipsoid centered at the cluster centroid. The radius of the ellipsoid along each dimension is the 90th percentile of the distance of all points in the cluster to the centroid along that dimension. Ellipsoids are colored according to the dominant spectral feature (Fig. 4; also see *Movie S1* for better 3D visualization). These ellipsoids are analogous to 3D error bars that help visualize the approximate location of the clusters in the PCA space.

we characterize the space of activity $X(t)$, in Fig. S2 we characterize the space of velocities approximated as $X(t + 1) - X(t)$. Taken together, the results in Fig. S2 and Fig. 2C imply that the space of activity is low dimensional, whereas the fluctuations essentially are multidimensional noise. This suggests that some locations in the activity space are stabilized.

Brain Activity During ROC Exhibits Clusters Consistent with Metastable Intermediate States. Brain activity during ROC does not evenly occupy the volume spanned by the first three PCs, as evidenced by distinct peaks in the probability distribution shown in Fig. 3B. Consistent with abrupt fluctuations in spectral power (Fig. 2B), the data projected onto the first three PCs form eight distinct clusters (SI Materials and Methods), the approximate locations of which are shown in Fig. 3C. Although clustering was performed on the data concatenated across all experiments, the distribution of data from each experiment taken individually also was multimodal (Fig. S5A and B). Furthermore, the concordance of clustering between individual experiments is statistically significant (SI Materials and Methods, Figs. S5 and S6). Thus, the clusters represent reproducible and distinct states distinguished by the distribution of spectral power across brain regions.

Three lines of evidence indicate that these clusters represent attractor states of the thalamocortical dynamics: (i) The transitions between states are abrupt (e.g., Fig. S3), and the paucity of points between the peaks of the probability distribution (Fig. 3B group data, Fig. S5 individual experiments) suggests that the system does not spend a significant amount of time between the densely occupied states. (ii) Dwell times within each state may last up to several minutes (Fig. S7A). (iii) Fluctuations die down when the system arrives into the clusters and increase between clusters (Fig. S8). The decrease in the amplitude of fluctuations associated with the arrival into densely populated regions of the parameter space suggests stabilization of neuronal activity. In this view, the multimodal distribution of brain activity in PCA

space may be seen as an anesthetic–concentration-dependent energy landscape in which the location of local energy minima gives rise to densely occupied states and local maxima demarcate boundaries between them. Note that the stabilization is not enough to trap the brain in any one state permanently, and spontaneous state transitions are observed readily at many anesthetic concentrations. Thus, we refer to the densely occupied regions of the parameter space as metastable states.

The characteristic spectral profile for each state (Fig. 4A) reveals that they can be grouped further into three distinct categories. Although each group of states exhibits a consistent increase in power at distinct frequency bands observed across all anatomical sites, individual members of each group are distinguished by the anatomical distribution of power in the high-frequency range. This suggests that fluctuations observed between clusters within the same group correspond to state-dependent fluctuations in thalamocortical coupling en route to awakening.

Clustering allows us to simplify ROC further as a sequence of states, starting from burst suppression and ultimately leading to wakefulness (Fig. 4B). The observed sequences of states reveal an additional element of the structure—some state transitions appear more frequently than others. Note, for instance, that although burst suppression is exhibited in all blue states, transitions to δ -dominant states (green) are observed from only one of them. Likewise, transitions to θ -dominant states (red) are observed from a single δ -dominant state. Although certain transitions appeared more frequently than others, to a good approximation the specific sequence of states appears stochastic. Thus, we assumed that the system is Markovian and computed transition probabilities between all pairs of states.

State Transition Probabilities Are Not Uniform; Rather, a Few States Are Hubs. Consistent with these observations, the network formed by state transitions linking the metastable states is highly structured—some states (blue and green asterisks) form hubs—arrival

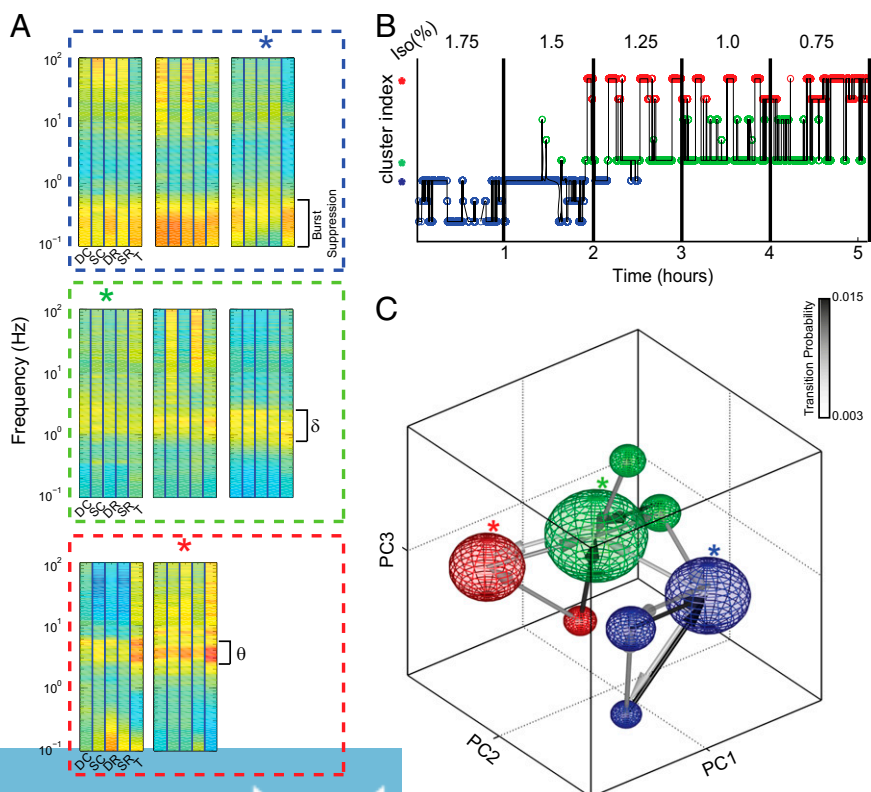


Fig. 4. The network linking the metastable states contains hubs; arrival into the hubs is essential for ROC. (A) Characteristic spectral density for each cluster (labels as in Fig. 3C). Clusters are grouped based upon the frequency range [blue, burst suppression; green, δ (1–4 Hz); red, θ (4–8 Hz)] that has consistently increased power in all electrodes. (B) The trajectory of a single experiment through the clusters suggests that certain state transitions are more likely than others. (C) A sphere centered at the cluster centroid shows each cluster. The radius of the sphere is proportional to the total time spent in the cluster. Arrow color shows transition probability. The network of transition probabilities reveals two hubs (blue and green asterisks in A–C) defined as the targets of multiple convergent transitions. Awakening was observed in the cluster shown by red asterisks ~95% of the time. Note that all paths into the “awake” cluster (red asterisks) from burst suppression (any one of the blue spheres) must involve passage first through the burst suppression hub state (blue asterisk) then through the δ -dominant hub (green asterisk). See also Movie S2 for better visualization of the network structure.

into which is a prerequisite for further progression to wakefulness, which was observed when the animal was in the cluster marked with the red asterisk ~95% of the time (Fig. 4C). This observation is confirmed by a formal analysis of centrality (*SI Materials and Methods* and Fig. S9). Note that by following the links in the network, the system starting out from any one of the burst suppression states must necessarily first arrive into the burst suppression hub (blue asterisk) and then the δ -dominant hub (green asterisk). Furthermore, although fluctuations between several δ -dominant states exist, transition to the awake cluster (red asterisk) occurs from the δ -dominant hub. Thus, the structure of the network that links the different metastable states gives rise to an orderly progression from burst suppression to consciousness.

It has been argued based on pharmacologic manipulations (21, 22), functional brain imaging (23), and deep brain stimulation (24) that activation of intralaminar thalamic structures is essential for ROC. Indeed, we find that the unifying characteristic of the θ -dominant states (red, Fig. 4A–C and Fig. 3C) is an increase in high-frequency oscillations in the intralaminar thalamus, signifying increased neuronal firing. Note, however, that although the two θ -dominant clusters exhibit similar degrees of thalamic activation, they are distinguished by the degree of cortical recruitment. Thus, although activation of the thalamus is an essential feature of brain activity associated with ROC, its presence is not sufficient to predict activation of the cortex.

The greater prevalence of hub states is attributable to more frequent arrivals into the hub rather than to longer dwell times (Fig. S7). This is consistent with convergent transitions into the hubs from multiple states. Together with the dimensionality reduction and clustering, the structure of the network that links different metastable states tremendously simplifies the problem of ROC after a drastic anesthetic-induced perturbation to brain activity.

Discussion

Here, we demonstrate that ROC involves discrete, individually stabilized patterns of neuronal activity. Furthermore, the network that links them gives rise to an orderly progression through these activity patterns toward eventual ROC. Because transitions toward the patterns of neuronal activity consistent with consciousness are observed from only a small subset of available states that we identify as hubs, arrival into the hubs may be used to suggest the possibility of impending ROC.

ROC from Isoflurane Anesthesia Is a Series of Transitions Between Discrete Metastable Intermediate States. It has been known since the 1930s that for most anesthetic agents, increasing depth of anesthesia correlates with lower-frequency, higher-power EEG oscillations until the onset of burst suppression (25). As a result, many depth-of-anesthesia measures, including multiple commercially available options, depend upon the frequency range that contains the dominant power. Previous work, however, was silent on whether this represents a continuum or discrete states. We show here that it is possible to detect a few discrete metastable intermediate states in multiple animals emerging from isoflurane anesthesia. Moreover, certain of these intermediate states function as hubs during ROC, such that transitions between groups of intermediate states necessitate transiting through the hub.

Certain aspects of the dynamics of ROC may in fact explain some of the difficulties faced by current depth-of-anesthesia monitors. For example, one consequence of the “stickiness” of the intermediates might be the relatively poor correlation between the depth-of-anesthesia monitor and the end-tidal anesthetic concentration in the B-Unaware trial (26). Another common feature when working with the available depth-of-anesthesia monitors is the occasional large jump with no apparent stimulation or change in anesthesia level, which might accord with a sudden state change with no precursor. Finally, several of the

current commercial systems output a number in a range from 0 to 100 to indicate depth of anesthesia; it is unclear whether the implied continuity of a 0–100 scale is optimal if ROC is characterized by a series of a few discrete states.

Generalizability of Particular Metastable Intermediate States. Although the consistency of clustering is statistically significant (Figs. S5 and S6) and suggests some common activity patterns observed across animals, there are clear differences in the distribution of activity among animals (e.g., Fig. S5). This variability may arise from biological factors, such as intrinsic differences between animals, that may reflect genetic, environmental, and developmental differences, as well as differences in the sensitivity to anesthetics. Other sources of variability may be experimental, such as differences in the precise location of the electrodes, as well as differences in the properties of the electrodes, such as impedance. Furthermore, some variability may be imposed onto the data by the analysis technique, such as truncation of the PCA after three PCs.

That being said, some of the observed variability may be a consequence of the stochastic nature of the state transitions. Indeed, although five of six animals visit each of the eight clusters, Fig. S10 shows that in some cases, a particular animal may contribute a disproportional amount of data to a cluster. One reason for this is that the transition probabilities differ somewhat among different animals. Given that transition probabilities are relatively rare events, full quantification of the variability among animals will require much larger dataset. However, another reason for the observed variability is that although transitions between clusters are rare (Fig. S10B), the probability of staying in a cluster is high. Thus, a small difference in the number of visits to a particular cluster may translate into a large difference in the total time spent in the cluster.

Anesthetic Inertia. ROC after anesthesia is not simply an issue of drug washout kinetics (27); rather, it is impeded by previously unknown neuronal processes that give rise to inertia—a phenomenon conserved across evolution and independent of the specific choice of anesthetic (28). The most compelling evidence for this inertia is that the dose–response curve for ROC is left-shifted with respect to that for induction of anesthesia (28). Although this hysteresis implies that ROC cannot be explained solely in terms of pharmacologic actions of the anesthetic agent—to predict whether one is anesthetized or awake, the internal state of the brain also must be known—it does not imply any specific class of neuronal mechanism. Although the metastable states we report here trap the brain in an unconscious state and may give rise to hysteresis, by significantly decreasing the number of available activity configurations, they allow eventual recovery.

Generalizability of ROC from Anesthesia. It is likely that recovery from other brain perturbations is similarly characterized by abrupt state transitions. For instance, it has been suggested that recovery from a brain injury resulting in a minimally conscious state is characterized by the emergence of abrupt state transitions evident on the EEG (29). It is well known that natural sleep is characterized by abrupt state transitions (e.g., ref. 9). Although there is overlap in terms of both the activity patterns and neuronal networks involved in sleep and anesthesia (e.g., refs. 2, 27, 29), the metastable configurations identified herein are fundamentally distinct from those observed in sleep. For instance, burst suppression never is observed naturally. Here, we focused on recovery from anesthesia because of the advantages afforded by the ability to parametrically control brain inactivation. This allowed us to directly relate the degree of brain inactivation to the energy landscape that defines the available neuronal activity patterns and the topology of transitions between them.

These results exemplify a general solution to the problem of recovering from a strong perturbation. Even relatively simple nonlinear

dynamical systems exhibit complex behaviors characterized by multiple attractor states and bifurcations (30). Thus, a strong perturbation may result in the system being permanently trapped in a different attractor. Biological systems capable of recovery, such as the brain, therefore must possess some self-tuning mechanisms by which the system traverses the state space to recover from a perturbation. Although the details of the metastable states encountered en route to ROC may differ as a function of the specific perturbation and species, the need for dimensionality reduction, together with stabilization of individual activity patterns, is almost unavoidable given the potential complexity of brain dynamics.

Finally, for the purposes of this paper, we have assumed that brain activity under anesthesia is Markovian. However, the brain is not a memoryless system, and it is entirely possible that ongoing brain activity under anesthesia may result in slow destabilization of the current state, eventually resulting in an abrupt transition. Understanding both the topology of neuronal dynamics underlying ROC and the mechanisms that drive transitions between these states offers a new avenue to explore the substrates necessary for consciousness and a new set of constraints for any mechanistic theory of consciousness.

Materials and Methods

Surgery and Animal Care. All use of laboratory animals was consistent with *Guide for the Care and Use of Laboratory Animals* and approved by The Rockefeller University's Institutional Animal Care and Use Committee (*SI Materials and Methods*).

- Lammi MH, Smith VH, Tate RL, Taylor CM (2005) The minimally conscious state and recovery potential: A follow-up study 2 to 5 years after traumatic brain injury. *Arch Phys Med Rehabil* 86(4):746–754.
- Brown EN, Lydic R, Schiff ND (2010) General anesthesia, sleep, and coma. *N Engl J Med* 363(27):2638–2650.
- Tuckwell HC (1989) *Stochastic Processes in the Neurosciences* (Society for Industrial and Applied Mathematics, Philadelphia).
- Levinthal C (1968) Are there pathways for protein folding? *J Chim Phys* 65(1):44–45.
- Schiff ND (2008) Central thalamic contributions to arousal regulation and neurological disorders of consciousness. *Ann N Y Acad Sci* 1129(1):105–118.
- Alkire MT, Hudetz AG, Tononi G (2008) Consciousness and anesthesia. *Science* 322(5903):876–880.
- Percival DB, Walden AT (1993) *Spectral Analysis for Physical Applications* (Cambridge Univ Press, Cambridge, UK).
- Destexhe A, Contreras D, Steriade M (1999) Spatiotemporal analysis of local field potentials and unit discharges in cat cerebral cortex during natural wake and sleep states. *J Neurosci* 19(11):4595–4608.
- Gervasoni D, et al. (2004) Global forebrain dynamics predict rat behavioral states and their transitions. *J Neurosci* 24(49):11137–11147.
- Stratmann G, et al. (2009) Increasing the duration of isoflurane anesthesia decreases the minimum alveolar anesthetic concentration in 7-day-old but not in 60-day-old rats. *Anesth Analg* 109(3):801–806.
- Swank RL, Watson CW (1949) Effects of barbiturates and ether on spontaneous electrical activity of dog brain. *J Neurophysiol* 12(2):137–160.
- Stecker MM, et al. (2001) Deep hypothermic circulatory arrest: I. Effects of cooling on electroencephalogram and evoked potentials. *Ann Thorac Surg* 71(1):14–21.
- Ohtahara S, Ohtsuka Y, Yamatogi Y, Oka E (1987) The early-infantile epileptic encephalopathy with suppression-burst: Developmental aspects. *Brain Dev* 9(4):371–376.
- Thömkle F, Brand A, Weilemann SL (2002) The temporal dynamics of postanoxic burst-suppression EEG. *J Clin Neurophysiol* 19(1):24–31.
- Steriade M, Nuñez A, Amzica F (1993) A novel slow (< 1 Hz) oscillation of neocortical neurons in vivo: Depolarizing and hyperpolarizing components. *J Neurosci* 13(8):3252–3265.
- Nunez PL (2006) *Electric Fields of the Brain: The Neurophysics of EEG* (Oxford Univ Press, New York).
- Steriade M, Gloor P, Llinás RR, Lopes de Silva FH, Mesulam MM (1990) Report of IFCN Committee on Basic Mechanisms. Basic mechanisms of cerebral rhythmic activities. *Electroencephalogr Clin Neurophysiol* 76(6):481–508.
- Franks NP (2008) General anaesthesia: From molecular targets to neuronal pathways of sleep and arousal. *Nat Rev Neurosci* 9(5):370–386.
- Doyle PG, Snell JL (2000) Random walks and electric networks. *Carus Mathematical Monographs* 22:1–120.
- Jones EG (1998) Viewpoint: The core and matrix of thalamic organization. *Neuroscience* 85(2):331–345.
- Alkire MT, McReynolds JR, Hahn EL, Trivedi AN (2007) Thalamic microinjection of nicotine reverses sevoflurane-induced loss of righting reflex in the rat. *Anesthesiology* 107(2):264–272.
- Lioudyno MI, et al. (2013) Shaker-related potassium channels in the central medial nucleus of the thalamus are important molecular targets for arousal suppression by volatile general anesthetics. *J Neurosci* 33(41):16310–16322.
- Långsjö JW, et al. (2012) Returning from oblivion: Imaging the neural core of consciousness. *J Neurosci* 32(14):4935–4943.
- Schiff ND, et al. (2007) Behavioural improvements with thalamic stimulation after severe traumatic brain injury. *Nature* 448(7153):600–603.
- Gibbs FA, Gibbs EL, Lennox WG (1937) Effect on the electro-encephalogram of certain drugs which influence nervous activity. *Arch Intern Med* 60(1):154–166.
- Whitlock E, et al. (2011) Relationship between bispectral index values and volatile anesthetic concentrations during the maintenance phase of anesthesia in the B-Unaware trial. *Anesthesiology* 115(6):1209–1218.
- Kelz MB, et al. (2008) An essential role for orexins in emergence from general anesthesia. *Proc Natl Acad Sci USA* 105(4):1309–1314.
- Friedman EB, et al. (2010) A conserved behavioral state barrier impedes transitions between anesthetic-induced unconsciousness and wakefulness: Evidence for neural inertia. *PLoS One* 5(7):e11903.
- Victor JD, Drover JD, Conte MM, Schiff ND (2011) Mean-field modeling of thalamo-cortical dynamics and a model-driven approach to EEG analysis. *Proc Natl Acad Sci USA* 108(Suppl 3):15631–15638.
- Strogatz SH (1994) *Nonlinear Dynamics and Chaos* (Perseus, Reading, MA).
- Clement EA, et al. (2008) Cyclic and sleep-like spontaneous alternations of brain state under urethane anaesthesia. *PLoS One* 3(4):e2004.

Data Acquisition. Wideband cortical and thalamic activity sampled at 40 kHz, with synchronous electrocardiography sampled at 1 kHz, was recorded to hard drive using a commercially available multichannel op-amp system (Plexon). Twenty-four simultaneous neuronal recordings were made with a custom linear microarray manufactured by Alpha Omega with three “prongs” separated by 1 mm. The two cortical prongs had eight leads spaced at 250 μm (roughly corresponding to the cortical layers) and a thalamic prong with eight leads spaced at 100 μm beginning 3.5 mm below the deepest cortical electrode. The diameter of all contacts was 25 μm . This array was inserted parallel to the anterior–posterior axis. The most posterior prong was inserted at 3.6 mm posterior and 1.2 mm lateral to the bregma and advanced along the dorsal–ventral axis (~5.2 mm) such that the thalamic electrodes spanned the centrolateral nucleus.

Spectral Analysis. All analysis was performed offline using custom-written software in MATLAB (MathWorks). LFPs were extracted using an acausal fourth-order Butterworth filter with a low pass frequency of 500 Hz to minimize phase distortion and down-sampled to 1 kHz. Evolving spectral activity was determined by using the Thomson multitaper method (31) using a sliding 60-s window with a window step of 1 s using 17 tapers, yielding a frequency resolution of 0.15 Hz. Two hundred forty-nine power estimates were computed for a frequency range between 0.15 and 300 Hz for each channel.

ACKNOWLEDGMENTS. This work was supported by a Foundation for Anesthesia Education and Research grant and K08 (GM106144-01) awarded to A.P. and an American Association of University Women International Fellowship awarded to D.P.C.