Personalized models of Disorders of Consciousness reveal complementary roles of connectivity and local parameters in diagnosis and prognosis

Lou Zonca[1], Anira Escrichs[1], Gustavo Patow[1,2], Dragana Manasova[3,4], Yonathan Sanz-Perl[1], Jitka Annen[5,6], Olivia Gosseries[5,6], Steven Laureys[5,7,8], Jacobo Diego Sitt[3], Gustavo Deco[1]

June 12, 2025

Abstract

The study of disorders of consciousness (DoC) is very complex because patients suffer from a wide variety of lesions, affected brain mechanisms, different severity of symptoms, and are unable to communicate. Combining neuroimaging data and mathematical modeling can help us quantify and better describe some of these alterations. The goal of this study is to provide a new analysis and modeling pipeline for fMRI data leading to new diagnosis and prognosis biomarkers at the individual patient level. To do so, we project patients' fMRI data into a low-dimension *latent-space*. We define the latent space's dimension as the smallest dimension able to maintain the complexity, non-linearities, and information carried by the data, according to different criteria that we detail in the first part. This dimensionality reduction procedure then allows us to build biologically inspired latent whole-brain models that can be calibrated at the single-patient level. In particular, we propose a new model inspired by the regulation of neuronal activity by astrocytes in the brain. This modeling procedure leads to two types of model-based biomarkers (MBBs) that provide novel insight at different levels: (1) the connectivity matrices bring us information about the severity of the patient's diagnosis, and, (2) the local node parameters correlate to the patient's etiology, age and prognosis. Altogether, this study offers a new data processing framework for resting-state fMRI which provides crucial information regarding DoC patients diagnosis and prognosis. Finally, this analysis pipeline could be applied to other neurological conditions.

Introduction

The question of consciousness, or lack thereof, is a complicated and controversial notion that has attracted the attention of theologists[1], philosophers[2], psychologists, and medical doctors[3] for centuries and remains under active investigation. In recent years, scientists and researchers have been investigating novel ways to characterize and quantify different states of consciousness. In

[1]Center for Brain and Cognition, University Pompeu Fabra, Barcelona, Spain
[2]Girona University, Spain
[3]Sorbonne Université, Institut du Cerveau - Paris Brain Institute - ICM, Inserm, CNRS, Paris, France
[4]Université Paris Cité, Paris, France
[5]Coma Science Group, GIGA-Consciousness, University of Liege, Belgium
[6]NeuroRehab & Consciousness Clinic, Neurology Department, University Hospital of Liège, Belgium
[7]Joint International Research Unit on Consciousness, CERVO Brain Research Centre, Laval University, Canada
[8]International Consciousness Science Institute, Hangzhou Normal University, Hangzhou, China

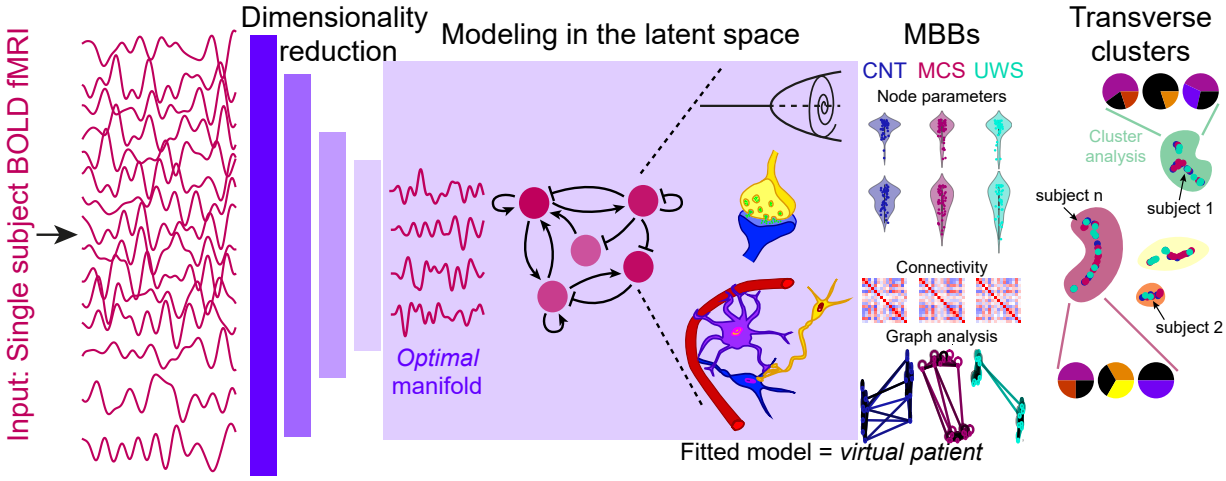

Figure 0: **Graphical abstract**: The BOLD fMRI time series (magenta) are projected into a latent space of low dimension (purple) using an auto-encoder (left). In the latent space, different latent whole-brain models are fitted to each subject's data (center). The fitted model parameters constitute 1st-level MBBs, allowing the differentiation between diagnosis (blue, magenta, and cyan). Finally, transverse clusters based on the fitted models can be related to the patient's etiology, age, and prognosis (right).

particular, disorders of consciousness (DoC) refer to a scope of conditions that can occur after patients who suffered various types of brain damage (e.g., traumatic brain injury (TBI), stroke, or cardiac arrest leading to anoxia) fall into a coma from which they only partially recover. There are two main dimensions used to characterize DoC patients: (1) *arousal*, which refers to wakefulness or vigilance, and (2) *awareness*, of the environment or the self[4]. Based on these two dimensions, medical doctors have defined several categories of DoC patients: the Vegetative State or Unresponsive Wakefulness Syndrome (VS/UWS)[5] corresponds to patients who do not show any conscious responses to sensory stimulation; the Minimally Conscious State (MCS)[6] describes patients who show limited but clear evidence of awareness. MCS has been further divided into MCS− and MCS+[7] based on the level of language understanding. However, each of these categories include very heterogeneous populations of patients in terms of pathology, behavior, and cognitive ability. The introduction and evolution of neuroimaging and electrophysiology allowed researchers to better quantify several aspects of DoC patients' responses, specifically the notion of covert consciousness, leading to a refined description of consciousness. This led to the discovery and definition of Cognitive-Motor Dissociation (CMD)[8], which describes patients who exhibit clear cognitive abilities despite the absence of motor responses. Using these techniques, consciousness has been theorized to be supported by cortical and subcortical awareness networks, involving multiple regions of the brain[4, 9, 10, 11, 12]. This led to the development of many quantitative scores or empirical biomarkers (EBMs) based on neuroimaging and/or electrophysiology data. Among them, the Functional Connectivity matrix (FC), which can be extracted from fMRI data at resting state (rs-fMRI), was proven to be sufficient to distinguish between UWS and MCS[13]. Furthermore, previous misdiagnosis of CMD patients[14] is a striking advocate for the development of accurate biomarkers. Indeed, a recent study showed that at least 25% of unresponsive patients could be classified as CMD patients [15]. Moreover, the chances of recovery for a CMD patient are significantly higher than for UWS patients[16] and a too early removal of life-support for such patients can have dramatic consequences.

However, although EBMs are theoretically driven, they remain purely descriptive of the patient's state and do not bring proper information about the underlying brain functions that are disrupted and therefore, cannot be used to investigate therapeutic approaches, which remain scarce today[17, 18, 11, 19]. In this context, the development of whole-brain computational models

emerged in parallel with the evolution of EBMs to bring a complementary understanding of the affected mechanisms in DoC patients. Indeed, for clinicians who aim to improve a patient's condition, it is fundamental to mechanistically understand what is missing in the patient's brain function to design a specific intervention that could restore the lost cognitive ability. These models are based on the fact that the dynamical co-activations and correlations between different small brain regions are shaped by the underlying structural connectivity (SC)[20, 21, 22]. Several studies have shown the crucial role of the SC and FC matrices in classifying patients into different groups representative of the patient's state using both fMRI[23, 24] and MEG[25] data during resting state. A correct classification of a patient's state is crucial for prognosis and decision-making, e.g., regarding the maintenance of life-support.

However, existing methods describe generic states in which patients can be categorized, but they drastically lack individuality as well as biological interpretability. Without individualization, there cannot be a systematic investigation of the underlying mechanistic or biological disruptions responsible for the patient's loss of consciousness, as they may be different from one patient to the next. Furthermore, these models solely explore the meso-scale inter-regional connectivity aspects underlying brain function, and none account for specific, local, biological processes that could play an important role in consciousness-related mechanisms. In particular, the regulation of neuronal activity by the astrocyte network could play a critical role in the diagnosis or recovery of some patients. Indeed, it is now well known that astrocytes play a crucial role in neuronal activity [26, 27, 28, 29, 30, 31, 32], which is also reflected in the BOLD signal [33] due to their strategic role in neurovascular coupling [34, 35, 36]. The main astrocyte pathways are metabolism regulation (through the astrocyte-neuron lactate shuttle), glutamate and GABA regulation, and extracellular potassium. Recent studies have shown the particular importance of the potassium pathway on neuronal bursting activity [37, 38, 39]. Other crucial reasons why we think that the regulatory role of astrocytes should be considered, in addition to its fundamental role in brain activity, are, first, that in severe brain lesions, astrocytes probably have been as impaired as neurons. Indeed, astrocytes are present in all brain regions and are even believed to be the most numerous cell type in the brain [40, 41]. Second, an increasing body of literature points to the fundamental role of astrocyte regulation of sleep, both in rapid eye motion (REM) and slow wave sleep[42, 43, 44, 45], which points to their active role in healthy consciousness mechanisms. Therefore, a solely neuron-centric approach seems limited if we want to understand the mechanisms underlying DoC, for at least some of our patients.

In this paper, we propose to investigate DoC patients, using biologically interpretable brain models that we will calibrate to each patient to grasp the differences between them. However, this approach requires working in very high-dimensional spaces, which leads to a prohibitive number of model parameters to fit and heavy computations. Furthermore, a large body of recent studies[46, 47, 48, 49, 50, 51, 52, 53] have unveiled the existence of a low-dimensional space that can encompass the information contained in the observed brain resting-state activity such as fMRI recordings. This reduced space must still account for the nonlinearities inherent to the neuronal activity that contains all the complexity underlying resting state brain dynamics and is commonly called the *latent space* or *latent manifold*.

In the present work, we started by investigating the optimal dimension to take the best advantage of our data without over-complexifying our models. In particular, we combined different dimensionality reduction methods (namely, auto-encoders and principal component analysis) to find the optimal (in a sense that we will define) latent space to project our data. Once the data was projected in the latent space, we built *latent* whole-brain models that we calibrated to each patient individually. Specifically, we compared a classical Hopf model with a new, biologically inspired model that accounts for astrocyte regulation of neuronal activity through the potassium pathway described above. Finally, we used the fitted model parameters as model-based biomarkers (MBBs), which we analyzed at two different scales: The fitted connectivity matrices, called the

generative effective connectivity matrices (GEC), showed that they carry the information about the severity of the patient's diagnosis. We then studied the connection between latent dimensions and the resting-state networks [54] in the natural space. Finally, we took a closer look at the local node parameters of the models and revealed that they carry information related to the patient's etiology, age and prognosis.

1 Results

1.1 A low-dimensional latent space preserving information

The main difficulty when fitting a model to data is having sufficient information to constrain the values of the, sometimes high, number of free parameters. Fitting meso-scale brain models, ranging from around 100 to 1000 nodes (depending on the initial atlas chosen) at the single-patient level thus constitutes a challenge. Indeed, the amount of information available in one fMRI recording of 5 min (approximately 300 time-points in our case), although informative, is limited, compared to the more classical approach of fitting one model for a group of dozens of subjects suffering the same condition. To overcome this limitation, we decided to explore data dimensionality reduction methods that would project the high-dimensional fMRI data (with as many dimensions as the parcellation size) into a lower dimensional space, also called the latent space. The reasons for doing this are multiple: first, our group recently showed that projecting the data in a reduced latent-space reduced the redundancies in the data, in addition to reducing the number of model parameters needed, thus allowing a more precise fit of the models [51]. Second, this approach also discards difficult open questions regarding fMRI preprocessing. Indeed, fMRI pre-processing is subject to various decisions on which there is still no consensus. These decisions include the choice of the parcellation atlas and whether to denoise or regress out several components of the signal (e.g., global signal regression). In this context, dimensionality reduction can be used to avoid these choices as we will look for the minimal dimension that preserves the information present in our data, thus eliminating any noise or redundancies in the signal.

In particular, we explored two open questions: (1) does the initial parcellation influence the dimension of the latent space and what is this *optimal* dimension? (2) Does the dimensionality reduction method matter?

1.1.1 The initial parcellation does not influence the dimension of the latent space

To determine the optimal dimension of the latent space, we trained Auto-Encoders (AE, Fig. 1A, purple) with different bottleneck dimensions ranging from 2 to 25. Briefly, an AE is an artificial neural network composed of successive layers of reducing dimension (the encoder) until reaching the bottleneck, or latent dimension, and a symmetric group of layers of increasing dimension (the decoder) that brings the data back to its original dimension. The AE is then trained to minimize the difference between input and output. The optimal latent space dimension is considered to be where the curve of the mean squared reconstruction error (MSE) vs latent dimension presents an elbow: indeed, when the latent dimension is too small, there is a loss of information and thus the reconstruction is bad, but then, increasing the latent dimension reduces the error until reaching the actual dimension needed to carry all the information contained in the data and, from this point forward, adding more dimensions does not improve the reconstruction (see methods 3.6.2 for more details).

Here, to determine whether the original parcellation influenced the results, we performed this analysis on the same data (see methods section 3.2 for description of the dataset) that we parcellated using the Schaefer parcellation [55] with 100 regions (Fig. 1B light purple), then, we added the subcortical regions using the Tian parcellation[56] with 16 regions (pink), and finally we tried the Schaefer parcellation with 1000 regions (deep purple). We found no difference in the MSE between

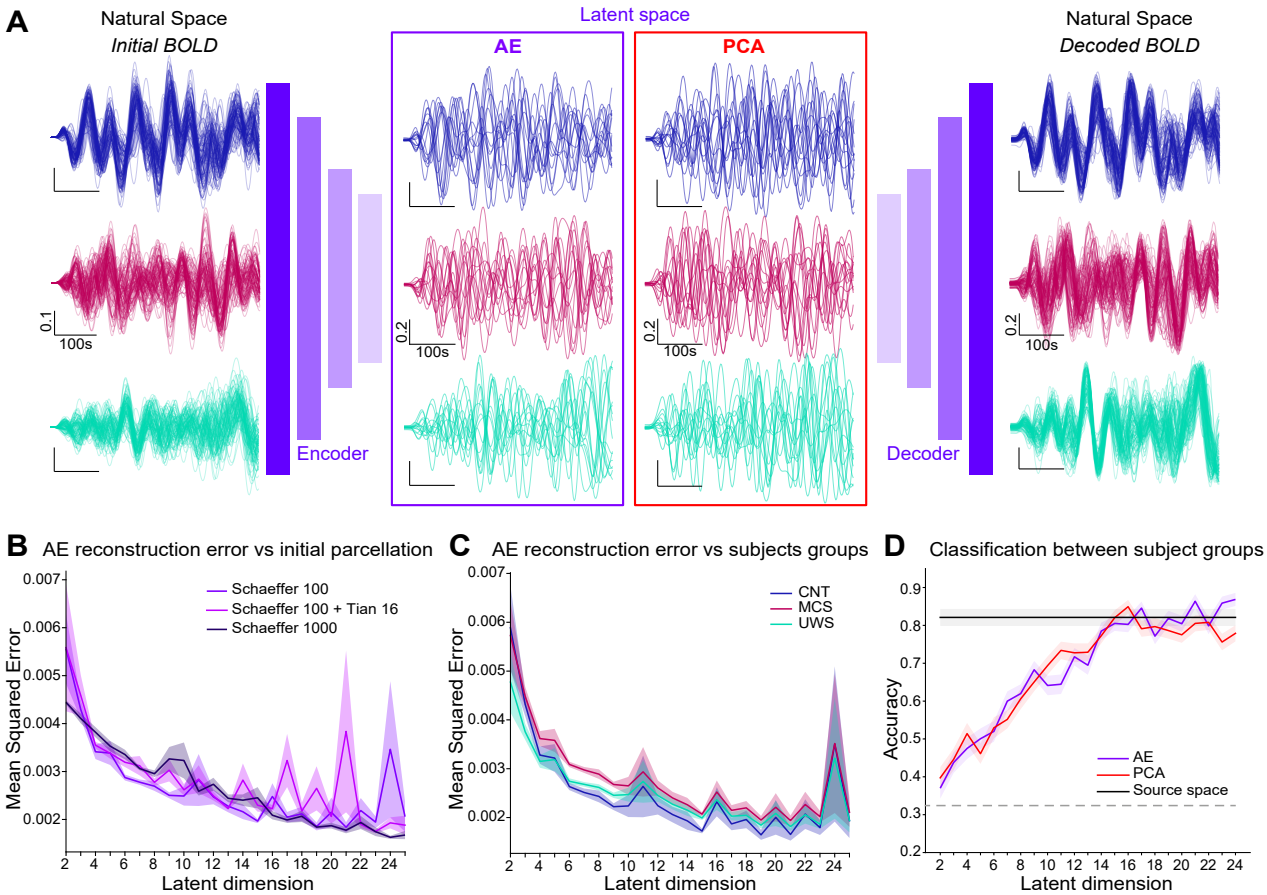

Figure 1: **Latent space exploration**: **A.** Example subjects' initial rs-BOLD time-series (left) color-coded per condition (blue: CNT, magenta: MCS, cyan: UWS) were projected into a reduced latent space (center) either using an AE (center left, purple) or PCA (center right, red) and then decoded (for the AE) trying to minimize the reconstruction error (right). **B.** MSE Reconstruction error for the AE with respect to the latent space dimension for different initial parcellations with and without sub-cortical regions (shades of purple). **C.** MSE vs latent dimension for initial parcellation Schaefer 100, divided per subjects condition. **D.** Comparison of the classification accuracy between conditions for AE (purple) vs PCA (red) for all latent dimensions, compared to accuracy in the natural space (black solid line). Grey dotted line indicates chance level.

the different parcellations. Furthermore, the elbow was reached in all cases around dimension 15. Therefore, for the rest of the results of this study, we used the Schaefer 100 parcellation. This was done using different datasets, from two different centers (in Paris and Liege) of fMRI data from DoC patients with UWS and MCS diagnosis and healthy controls (see Methods sections 3.2 and 3.6.1 for details).

We also checked whether the patient's diagnosis would influence the elbow's position by using the previously trained AEs (trained on a balanced mix of subjects from all diagnosis and healthy controls (CNT)). We then split the subjects of the test set by condition (i.e. patients' diagnosis: UWS or MCS, or healthy controls) and looked at the MSE by group vs the latent-dimension (Fig. 1C. control (CNT), blue, MCS, magenta, UWS, cyan), and found that there was no difference in the optimal latent-dimension based on the subjects' condition.

1.1.2 The latent dimension minimizes the MSE and allows the classification between conditions better than in the natural space

To make sure that the dimensionality reduction did not lead to a loss of information, we decided to use a second criterion: given that we already knew that, in the natural space we can distinguish between clinical conditions based on the FC matrices of the subjects [13], we wanted to make sure that, in the latent space we were still able to do it, at least as well. To do so, we trained a support vector machine (SVM) classifier in the source space (of dimension 100) to distinguish between the empirical functional FC of all our subjects that are grouped into three categories based on the clinical exams: UWS, MCS, and CNT. We could obtain a very good accuracy of over 82%. We then used a similar training procedure on the reduced data for all latent dimensions between 2 and 25 and for the two-dimensionality reduction methods (PCA and AE). The training of the SVMs in the reduced dimensions was done on the latent empirical FCs, i.e. the Pearson correlation matrices of the reduced signals, i.e. for each reduced dimension $n \in [\![2, 25]\!]$ we obtain latent FC matrices of dimension $n \times n$ on which we train the SVM classifier. The rationale was that the latent dimension should be the smallest one that allows for similar, or even better, classification accuracy than in the source space. We found that dimension $d_{opt} = 15$ was the optimum for both methods. The fact that both methods gave the same results is a good sign of robustness, showing that the intrinsic level of information contained in the data is of this optimal dimension, regardless of the dimensionality reduction method used.

To conclude, we found that both AE and PCA give similar results regarding the classification of the FC matrices between clinical conditions. However, PCA is a linear process while we know that the processes underlying the generation of the fMRI signals that we observe are intrinsically non-linear. For this reason we decided to pursue our study using the AE. Finally, we projected all the available fMRI data into this latent space of optimal dimension 15, to start working on the modeling step that we describe in the next section.

1.2 Modeling rs-fMRI activity in the latent space

In this section, we introduce two models that we used to deepen our understanding of the data. Our goal here was to use the Hopf model[57] (see Section 3.7.1) as a *control* model, in the sense that this model has already proven very valuable for fMRI analysis, and then to push forward our analysis by introducing a novel model, that we shall describe below.

As mentioned, we implemented two different models: (a) the Hopf model, and (b) a new model based on synaptic short-term plasticity that accounts for the regulatory effect of astrocytes on neuronal activity. This is done through the modulation of a neuronal regulatory mechanism called afterhyperpolarization (AHP) [58, 37]. This last model simulates the neuronal response, which we then combined with the Balloon-Windkessel model[59, 60] to simulate the final BOLD response. We decided to add this new model (that we refer to as the AHP model in what follows) to our study to explore whether this regulatory effect is affected depending on the conditions because of the known impact of astrocytes in the regulation of sleep and brain oscillations [42, 30]. Indeed, this model allows us to have a direct biological interpretation of the fitted parameters. In other words, it provides model-based biomarkers (MBBs) with a direct biological meaning.

1.2.1 Individual patient fitting procedure

For both models, there are two types of parameters: (i) the local, or node parameters, e.g., the bifurcation parameter (often denoted a) for the Hopf model, which determines the type of activity (random fluctuations or synchronized oscillations); or the parameters relative to AHP in the case of the second model. (ii) The connectivity matrix between the nodes, also known as the Effective Connectivity (EC) matrix, which gives the actual level of influence each node of the model has on

the remaining ones.

We implemented the model fitting procedure (Fig.2A) as follows: (1) We first fitted the node parameters by minimizing the distance between the empirical and simulated FC, i.e., the 15×15 Pearson correlation matrices between the time series of the 15 nodes. We then computed the structural similarity index (SSIM)[61] between the empirical and simulated FC, which we maximized as a fitting criterion (a SSIM of 1 means perfect similarity and 0 none). Indeed, previous studies have shown that the SSIM offers a good trade-off between absolute (e.g MSE) and relative (e.g. correlation) assessment of the differences between two matrices[62]. The parameter space was then sampled randomly until a threshold $SSIM \geq 0.9$ was reached, or until the maximum number of simulations N_{max} was reached. We kept the parameter set that maximized the SSIM. (2) The second step of the procedure was to iteratively fit the generative effective connectivity (GEC) for each patient by comparing the empirical FC to the simulated one in a similar manner as described above, using the previously fitted node parameters (see Methods, Section 3.8, for more details).

We found that the fit quality plateaued around 0.6 for the Hopf model (Fig. 2B, left) and was constant over iterations. Note that in [62], a SSIM of 0.4 is already considered as a good fit, here, we considerably increased it. In contrast, the SSIM of the AHP model (Fig. 2B, right) increased over the first iterations before stabilizing around 0.80-0.86 (depending on the group, the UWS had the best-fit quality, although the difference was small).

Since we have more node parameters (6 vs 2) in the case of the AHP model compared to the Hopf model, it is of no surprise that the fit would be of better quality. Therefore, in the following sections, we present the results obtained from this model. However, all analyses were performed on both models, with the results for the Hopf model being available in the supplementary information (SI). We used this double analysis as a control to see whether we get coherent results with both models to make sure that we did not see model artifacts but robust results. In brief, the AHP model provided us with a more precise fit and interpretation than the Hopf model, but both models point to similar results, as we shall see in the next sections.

1.2.2 Node fitting results: model-based biomarkers

Each patient's set of fitted node parameters constitutes model-based biomarkers (MBB, Fig. 2C). For the AHP model, we obtained a vector of dimension 6 for each patient (the fitted node parameters are, from top left to bottom right: the noise intensity σ, the facilitation resting state X_{eq}, the depression resting state Y_{eq}, and the three AHP parameters: hyperpolarization depth T_h, and medium and slow timerates $1/\tau_{mAHP}$ and $1/\tau_{sAHP}$). We first checked whether these parameters revealed differences between the three groups (CNT in blue, MCS in magenta, and UWS in cyan) but, apart from the slow AHP time rate $1/\tau_{sAHP}$ that was significantly increased (i.e. indicating a shorter AHP duration) in the MCS group compared to the controls, as well as a close to significant difference (p=0.07) between MCS and UWS, the other parameters didn't show any clear difference based on the clinical state. We found the same results for the Hopf model MBBs, which are only 2 (noise intensity σ and the bifurcation parameter a, see Fig. S1). However, each distribution was quite spread out thus revealing a wide within-group variability. At this point, we concluded that the node MBBs do not carry much information about the clinical state, and therefore do not allow to distinguish between conditions but that they may hold other type of information about the patients' conditions such as their etiology, age, underlying affected mechanisms, and other clinical variables, as we explore in the last section of this study. But first, we still need to recover and explain the difference between the conditions we know we can distinguish from (Fig. 1D).

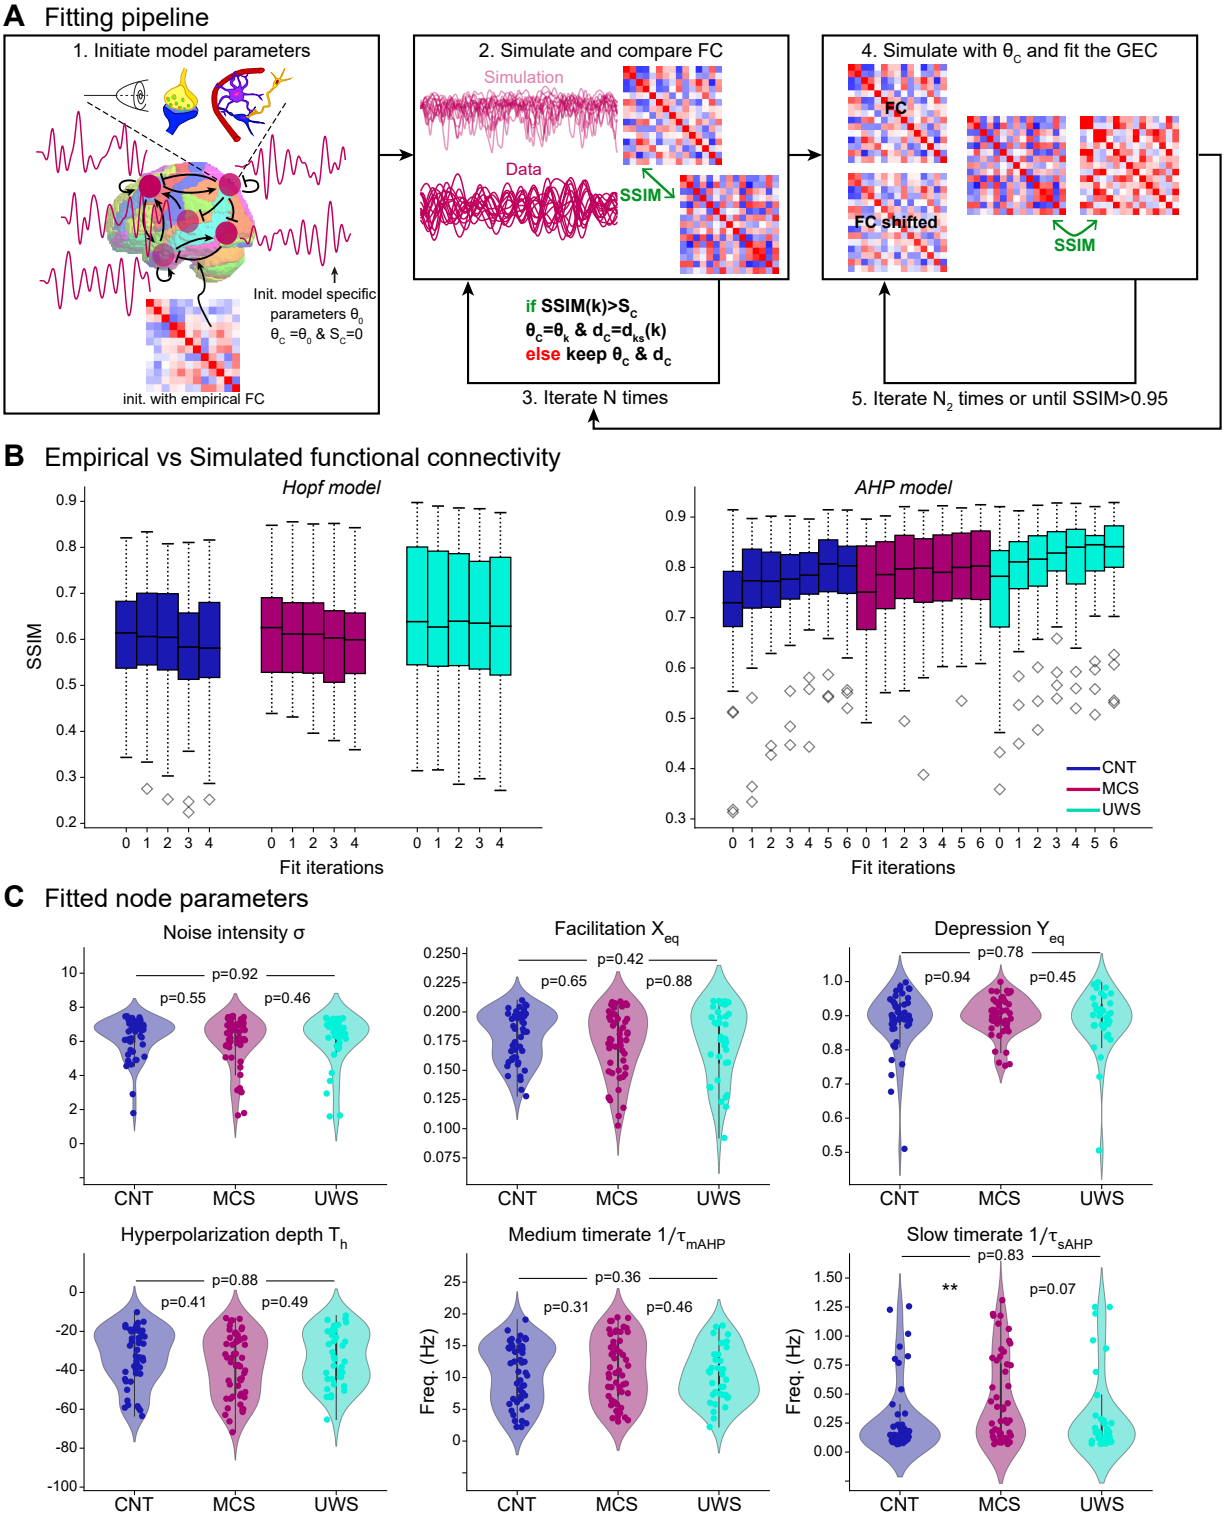

Figure 2: **Models fitting**: **A.** Iterative fitting procedure: the node parameters were randomly initialized and the connectivity was initialized by the empirical FC (step 1, left). Then, the subject's latent time series were simulated and the empirical and simulated FC were compared to maximize the SSIM (step 2, center) over the fitting iterations with random exploration of the parameter space (step 3). Then, with the best node parameters fixed, the GEC was fitted iteratively to maximize the SSIM between simulated and empirical FC and shifted FC (steps 4 and 5, right). The procedure was iterated at step 2, to fit the node parameters using the fitted GEC. **B.** Fitting results: SSIM between the simulated and empirical FC after various fitting iterations per condition (blue: CNT, magenta: MCS, cyan: UWS) for the Hopt (left) and AHP (right) models. **C.** 1st level MBBs per condition, obtained from the fitted node parameters (each dot is one subject, top left to bottom right: noise intensity σ, facilitation resting-state value X_{eq}, depression resting-state value Y_{eq}, hyperpolarization depth T_h, all dimensionless, medium $1/\tau_{mAHP}$ and slow $1/\tau_{sAHP}$ AHP time rates (in Hz), each dot is one patient.

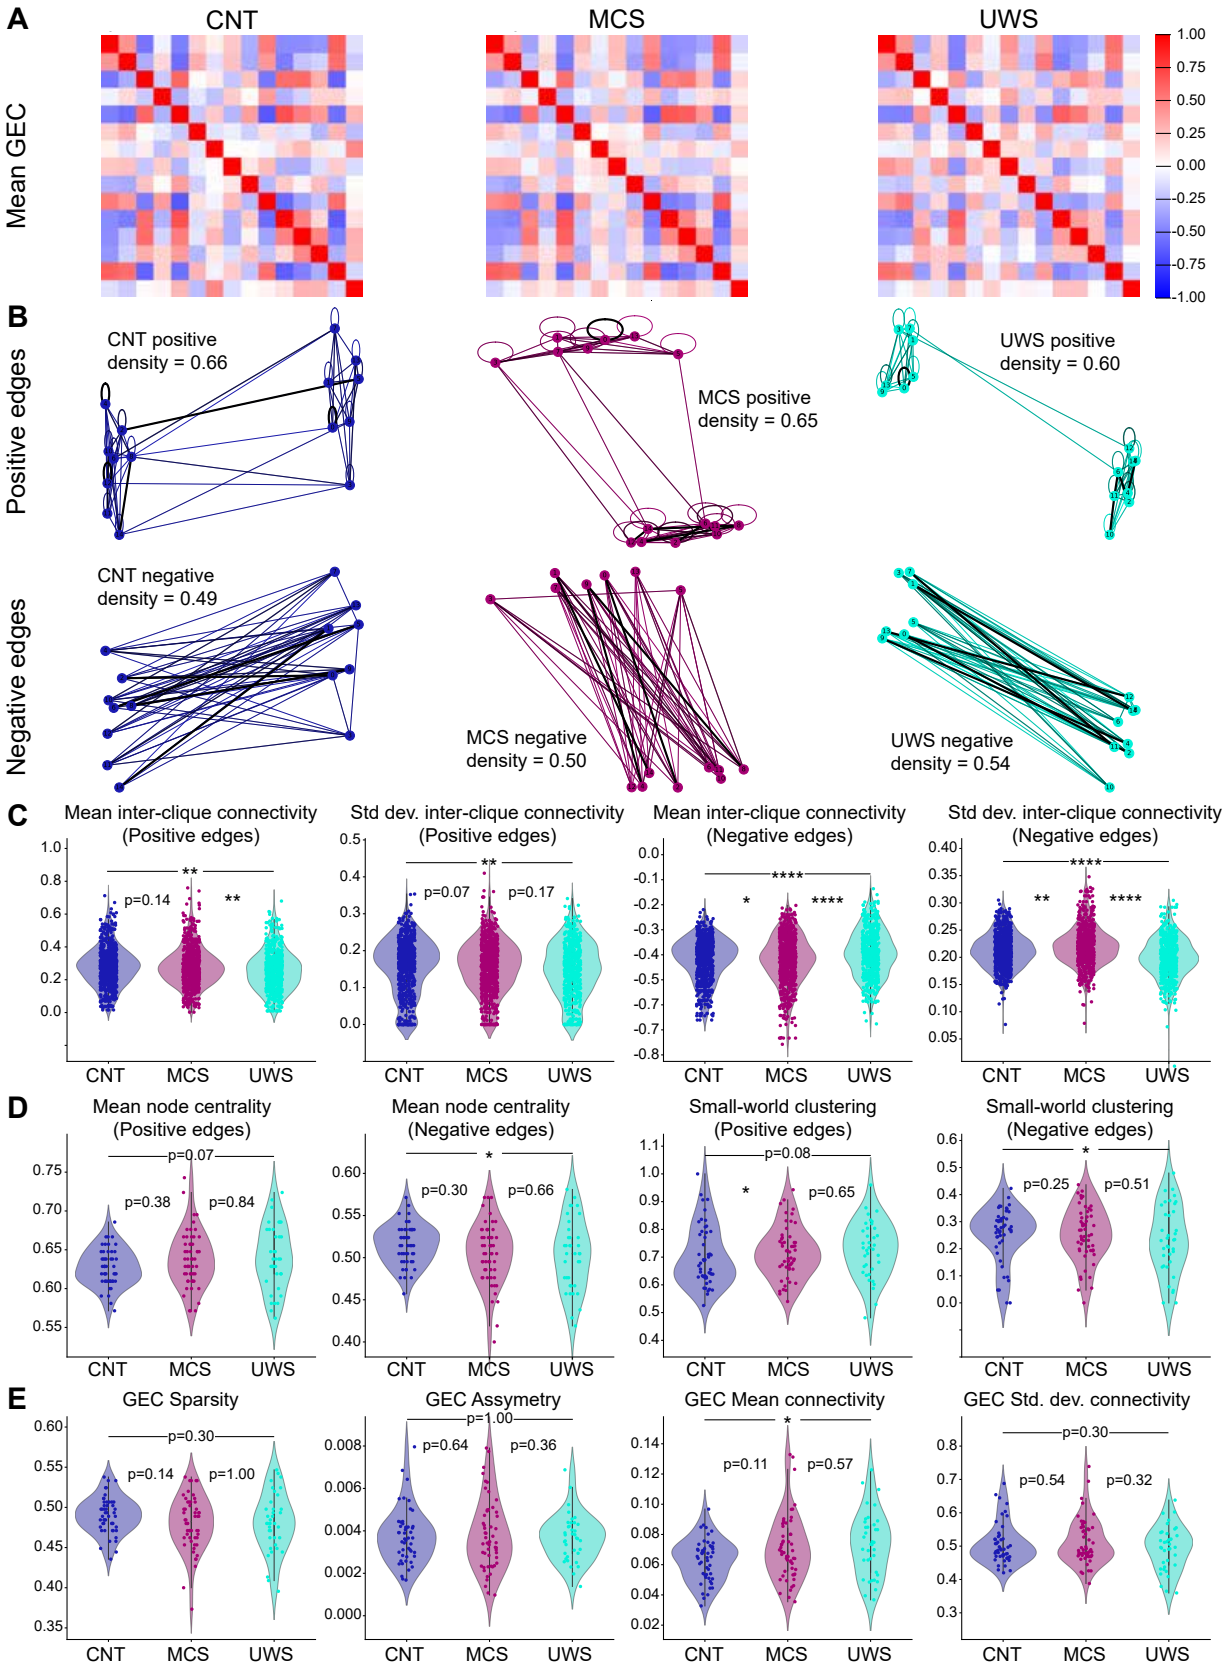

Figure 3: **GEC-based biomarkers for the AHP model**: **A.** Average over all subjects of the fitted GEC matrix per condition (left to right, CNT, MCS, and UWS). **B.** Extracted graphs from the average GEC matrices (top: positive connections, bottom: negative connections) showing two subgroups (or cliques) of nodes positively connected between them and projecting negatively on the other subgroup, for each condition. bf C. Inter-clique metrics: mean and standard deviation for positive (left) and negative (right) edges. **D.** Graph metrics: node centrality and Small-world clustering parameter. E. GEC extracted measures: left to right: sparsity, asymmetry, mean and standard deviation values. Statistical tests: pairwise Kolmogorov-Smirnov tests, ****: $p \leq 0.0001$, ***:$0.0001 < p \leq 0.001$, **:$0.001 < p \leq 0.01$, *:$0.01 < p \leq 0.05$

1.3 Analysis of the GEC reveals the differences between conditions

1.3.1 A two-clique structure for the GEC

To analyze the difference between the groups we thus turned to the GEC matrices. Fig. 3A shows the averaged fitted GEC matrix per condition. However, this visualization makes it hard to see the differences between them. Thus, to further explore the structure of these GEC matrices, we converted each GEC matrix into a weighted graph, where the GEC is the connectivity matrix of the graph. Then, for a clearer visualization, we split these graphs into two subgraphs: one with positive edges (Fig. 3B, upper line) and another one with negative edges (lower line). The first striking observation was that, for all groups, these graphs are constituted of two groups of nodes (or cliques), strongly positively connected within each group and negatively projecting to the other clique. At first glance, we seem to have a stronger within-clique positive density for the CNT (0.66), than for the MCS (0.65) and lowest for the UWS (0.60), while we have the opposite trend for the negative inter-clique projections (CNT: 0.49, MCS: 0.50 and UWS: 0.54).

We thus decided to further quantify the clique interactions (Fig. 3C). To do so, we used the *networkx* package to automatically compute the two major cliques based on the selection of each node, and computed the mean and standard deviation of the inter-clique connectivity for the positive (Fig. 3C, left), and negative subgraphs (right). We found that these measures, not only allowed us to differentiate the controls from the patients, but also between the two patients groups with high significance.

We also looked at other graph measurements (Fig. 3D), namely the mean node centrality for the positive subgraph (first left), which displayed a small difference (p=0.07) between CNT and MCS but not between the others; and the mean node centrality of the negative subgraph (second left), which allowed to distinguish between CNT and MCS but not between the others. The small-world clustering (Fig. 3D, left) for the positive edges showed a significant difference between CNT and MCS and a fairly significant difference (p=0.08) between CNT and UWS, but failed to distinguish between MCS and UWS, and the same for the negative edges, where it could only differentiate between CNT and UWS.

Finally we also directly quantified values from the GEC matrices (Fig. 3E): from left to right: the sparsity (quantified as the proportion of connections < 0.015), the asymmetry (difference between the GEC matrix and its transpose), and the mean connectivity and its standard deviation (i.e., the mean and std values of the matrix components). Similarly, we found no clear difference between the two DoC groups. However, the mean connectivity seemed higher in both DoC groups than for the CNT, indicating a more synchronized activity, as expected. We saw the same trends using the Hopf model (see Fig. S2).

In conclusion, it seems that the difference between conditions is mostly carried by this two-clique structure we observe in the GEC graphs, with stronger negative inter-clique and weaker positive inter-clique projections in the least conscious patients than in the more conscious ones and the controls.

1.3.2 Decoding the latent dimensions correlates to the resting state networks

To further understand the meaning of the two-clique structure observed in the GEC matrices (Fig. 3B), we decoded the signals from the latent GEC cliques and each latent dimension to see how they correlate to the natural structures in the brain. In particular, we investigated the link between GEC cliques (resp. latent dimensions) and the 7 Yeo Resting State Networks (RSN) [54], with a similar approach to the one used in [51] (see Methods, section 3.10).
Briefly, we decoded signals with activity (from the control subjects) on only one of the two GEC cliques. We computed the resulting FC matrix (Fig. 4A, left, and B) for the average over all control subjects. We then compared these two FC matrices with those corresponding to activity on each of the 7 RSNs (Fig. 4A, right). For both cliques, we computed the SSIM between the clique

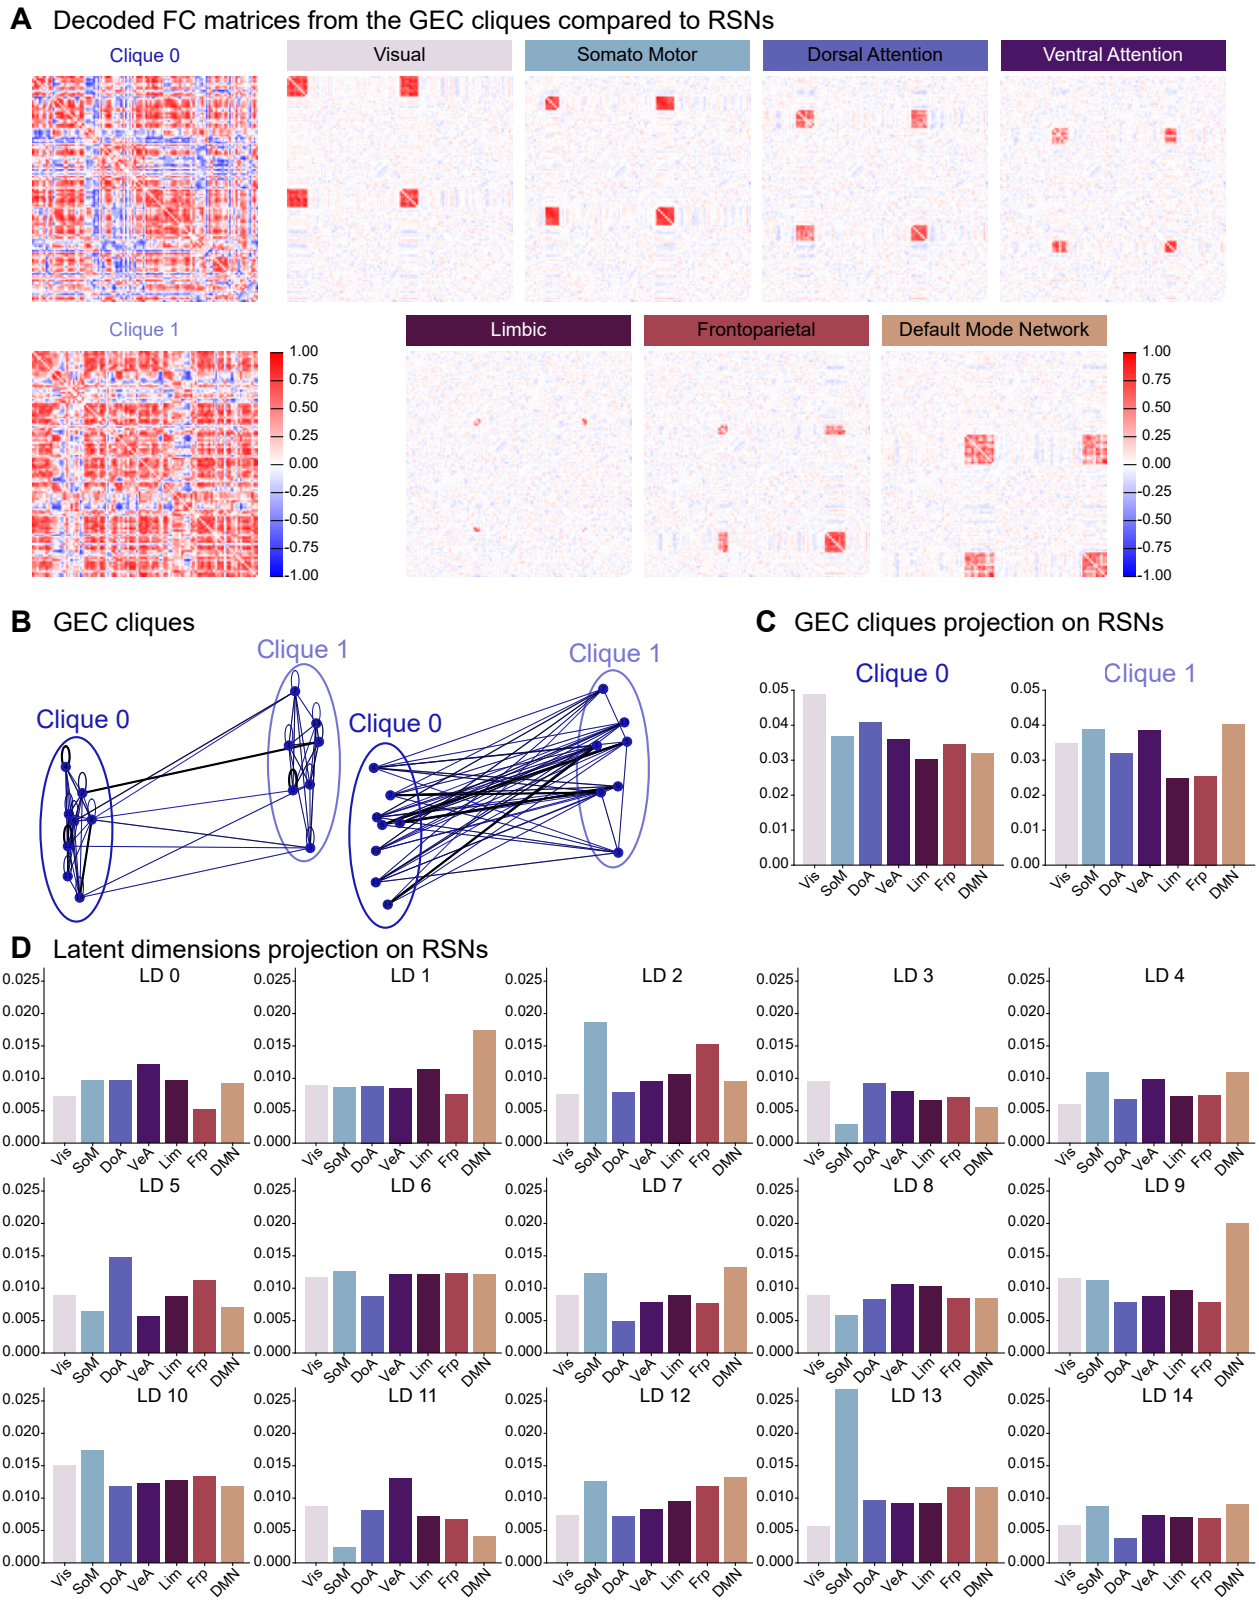

Figure 4: **GEC cliques and latent dimensions connection to the Resting State Networks**: **A.** FC matrices from the decoded signals for both GEC cliques (left), and FC matrices corresponding to the 7 RSNs in the natural space. **B.** GEC cliques from the average CNT GEC graphs. **C.** SSIM between the decoded FC matrices from each GEC clique compared to each RSN. **D.** Same as C but for the decoded signals from each latent dimension.

FC and each RSN FC (Fig. 4C). We found that the visual network (light gray) projected most strongly on Clique 0, followed by the Dorsal Attention network (blue). In contrast the default mode network (DMN, beige) and the Ventral Attention network (purple) were more strongly associated with Clique 1.
We then decided to investigate the connection between each latent dimension $LD \in [|0, 14|]$ and the 7 RSNs. Using a similar approach, we compared the FC matrices from decoded signals with activity on only one of the LD at a time to the RSN FC matrices (Fig. 4D). We found that the Somatomotor network (light blue), was most strongly associated with latent dimensions $LD13$ and 2, while the DMN strongly projected onto $LD1$ and 9. We also noted that $LD11$ was mainly related to the Ventral attention network. Generally, most networks were associated with a combination of latent dimensions rather than one in particular.
This could mean that the lower within-clique density observed in UWS patients is related to lower or less consistent activity on specific RSNs. However, a more detailed analysis would be required to confirm such a hypothesis.

1.4 The node MBBs reveal transverse clusters that correlate to etiology and outcome

In this section, we used a subset of our patients for whom we had more detailed diagnosis (splitting MCS between MCS+, and MCS−) as well as detailed metadata, including age, gender, etiology and evolution of the condition (different Coma Recovery Scale-Revised (CRS-R) diagnosis at different points in time after the patient's admission) as well as the outcome (see methods section 3.5).

As we pointed out before, DoC patients are a very heterogeneous group in symptoms and causes. Thus, we decided to explore the parameter space (the virtual patients or node MBBs) while being agnostic on the diagnosis. The first step was to look for a low-dimensional representation of the parameter space that could allow us to visualize the emergence of transverse clusters of patients for both models. We did so by using Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP)[63] (Fig. S3). We then recovered the emerging clusters that we saw in Fig. S3A using the *HDBScan* clustering algorithm[64] (Fig. S3B, 5A and S4A).

We could then see which parameters carried the specific distinction between clusters (Fig. 5B). For example, separation into cluster 0 (red) is mainly carried by the facilitation resting state parameter X_{eq} (Fig. 5B top, center, red), as well as the hyperpolarization depth T_h (bottom left), while the distinction between clusters 3 (light green) and 4 (purple) seems driven by a combination of all three AHP parameters (bottom row).

To understand the interpretation of these clusters, we checked how belonging to a specific cluster correlated with the available metadata (Fig. 5C). We plotted the age distribution of subjects of each cluster (Fig. 5C. top left), and the gender balance (top right), neither of which exhibited clear tendencies. Then the etiology (anoxia, traumatic brain injury (TBI), combined anoxia and TBI, stroke, SAH, and all others, Fig. 5C bottom left) and the outcome (see Methods section 3.11.1), which we split into dead, decreasing (i.e., when the evolution of the CRS-R diagnosis went towards less conscious over checks but did not lead to death), stable, increasing (when the diagnosis ameliorated without leading to consciousness), and conscious/emerged (Fig. 5C bottom right). Here we saw some very interesting tendencies: for example, clusters 2 and 3 contained mainly patients with negative outcomes, with a majority of deaths (respectively: 66.7% and 60%) and some stable DoC patients (33.3% and 40%) and 0 recovery in either cluster, while half of cluster 1 consists of emerged patients (50%). Finally, clusters 0 and 4 contained more mixed outcomes including dead (resp. 50% and 31.25%), emerged (25% each), and DoC patients (resp. 25% and 43.75%), as summarized in Table 1.

Since it is known that those factors (in particular, age, etiology, and outcome) are not independent (for example, we know that anoxia patients have a poorer prognosis), we further analyzed the composition of the clusters (Fig. 5D): for each cluster, we split the patients between emerged, dead

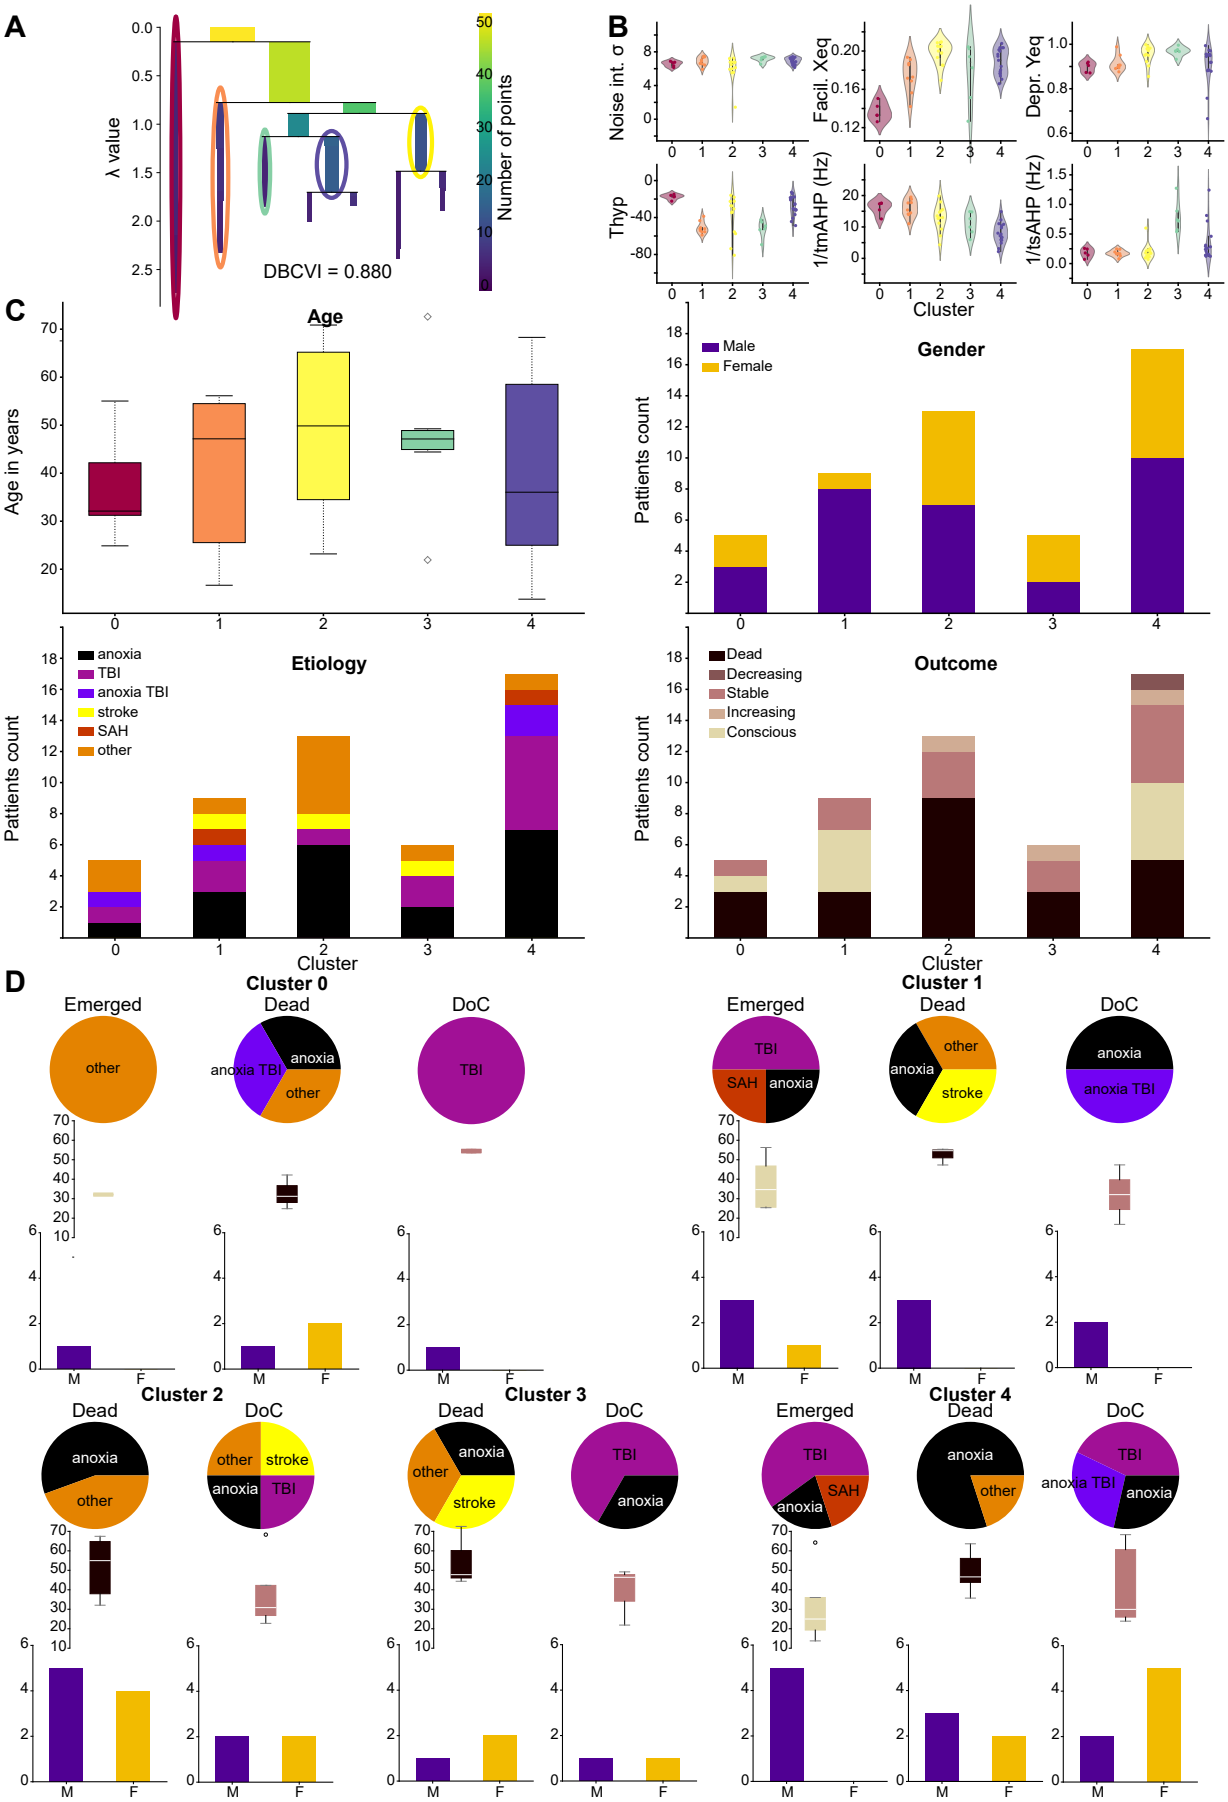

Figure 5: **Transverse clusters from the AHP model node parameters**: **A.** Condensed tree plot of the *HDBSCAN* clustering algorithm showing the level of similarity between clusters. **B.** Fitted node parameters, separated by cluster, showing which parameters influence the attribution to which clusters. **C.** Composition of the clusters: top left, the age distribution in years of the patients belonging to each cluster. Top right, gender distribution within clusters. Bottom left, etiology, bottom right, outcome. **D.** Co-factor analysis within clusters: for each cluster, we split the patients between emerged, still in DoC, and dead, and show the distribution of their etiology, age, and gender balance.

and DoC (grouping, increasing, stable and decreasing patients) and then looked at the distribution of etiology, age and gender with each category. Here, we see that for the clusters with mainly negative outcomes (clusters 2 and 3) both age and etiology correlate to the prognosis, with all the dead being older with 53, resp. 55, years old on average and who mainly suffered from anoxia (5 of the 9 deceased patients in cluster 2), while the stable DoC patients were the younger ones (39 years old on average in both clusters) and all of the ones with TBI associated to these clusters. However, compared to other TBI DoC patients who were associated with clusters with more favorable outcomes, this could be interpreted as a potentially worse prognosis since their MBBs were closer to those of people with a negative prognosis. In the same way, when we looked at cluster 1, which contains mainly favorable outcomes we saw that all the TBI patients from this cluster emerged, as well as some anoxic and the SAH ones. In the meantime, we saw once again that the few dead patients in this cluster (3 subjects) were the oldest ones (52 years old on average). When we looked into the mixed outcome clusters (0 and 4), similarly, we saw that age is a good predictor of the outcome since the emerged ones were the youngest patients (32, years old on average in both clusters), while the dead ones were almost all anoxic and older patients (33, resp. 49 years old). Finally, we note that gender did not seem to exhibit any clear tendencies regarding its correlation to one or the other of the observed co-factors (see Table 1).

Going back to the distribution of the node MBBs by cluster (Fig. 5B.), we could see that, for the negative outcome clusters 2 and 3 (yellow and light green) facilitation was quite high (similarly as in cluster 4, but higher than in the good outcome cluster 1). Facilitation can be associated with pre-synaptic calcium concentration, this difference could thus hint at a deregulated calcium concentration. Still, for clusters 2 and 3 we saw that depression was close to one, which can be interpreted as unaltered compared to healthy people (a depression resting state parameter of 1 means that the readily releasable pool of pre-synaptic vesicles is filled at rest). However, we could see that cluster 3 (green) showed a higher slow AHP time rate, which corresponds to a shorter AHP period. In previous studies [37, 58] we have associated this reduction in AHP with a disrupted astrocyte network and reduced astrocyte regulation of extracellular potassium levels. Similarly, we saw that cluster 2 seemed to be associated with a less deep hyperpolarization (parameter T_h). On the other hand, we have seen that facilitation (associated with presynaptic calcium) was very low in cluster 0 and the middle in cluster 1, while the distribution for cluster 4 seemed bimodal, with one-half at the level of cluster 1 and the other half closer to the values for clusters 2 and 3. Finally, cluster 0 seemed to have also a lower depression parameter, indicating more depressed synapses than in the healthy case. Finally, we summarized the results of the clusters composition in Table 1.

	Emerged	Dead	DoC	M	F	Age	TBI	Anoxia	Other
Cluster 0	25%	50%	25%	25%	75%	37.08 ± 10.55	25%	25%	50%
Cluster 1	50%	37.5%	12.5%	87.5%	12.5%	41.23 ± 14.05	25%	37.5%	37.5%
Cluster 2	0%	66.7%	33.3%	50%	50%	48.48 ± 16.09	8.3%	41.7%	50%
Cluster 3	0%	60%	40%	60%	40%	47.08 ± 14.70	20%	40%	20%
Cluster 4	25%	31.25%	43.75%	56.25%	43.15%	41.06 ± 17.60	31.25%	56.25%	12.5%

Table 1: Summary of clusters composition

In conclusion, we saw that these novel node MBBs bring us two new types of information: (1) based on these MBBs, we can group the patients, regardless of their current diagnosis, into clusters that are correlated with the outcome and etiology, thus helping for a future, more precise diagnosis and prognosis. (2) The direct biological interpretation of these parameters can help us investigate the underlying affected mechanisms and form new hypotheses regarding the patient's condition and their potential treatment.

2 Discussion

Latent space dynamics and structural interpretation A growing body of literature shows evidence of a low-dimensional manifold that encompasses the brain's resting-state activity. The dimension of this latent space is always found to be of the same order of magnitude: around 10 to 15, across various studies, both from our group and from others working on different datasets and conditions [50, 49, 48, 51, 53, 52]. Here, using rs-fMRI recording from DoC patients, we confirmed that finding and showed that we could work in a latent space of dimension 15 without losing the information contained in our data, or even our ability to explain it. This latent space description is not only convenient regarding the actual working dimension for explaining the data, but it also allows us to avoid inherent difficulties linked to fMRI preprocessing, parcellation choices, and denoising, as we showed in Fig. 1. This reveals that the intrinsic amount of information contained in rs-fMRI data, from two different datasets, can be explained in such tractable dimensions, regardless of the condition studied.

However, the question of the biological interpretation of the latent space remains open. We believe that it should be interpreted in terms of interacting networks. In a recent study, we showed how the dimensions of the latent space (also called latent modes) project onto the resting state networks (RSNs)[51]. Here, we deepened this analysis of the connection between the latent dimensions and the RSNs (fig. 4). We found that each RSN seemed to be encoded by a combination of latent-dimensions, confirming that the results observed in the latent space have a biological interpretation. However, given the complexity and the amount of possible combinations, further research is still needed to decipher fully the mapping between the latent and natural spaces. Meanwhile, other works are unveiling the tight relation between spatial and temporal activity in the brain, or in other terms, how structure and function are intrinsically linked. For example, using harmonic decomposition, Atasoy and colleagues showed in several studies how the RSNs overlapped with functional task-based networks [65], thus revealing harmonic brain modes that could be compared to the latent modes. This framework recently allowed us to reveal and quantify the effect of hallucinogenic drugs such as DMT [66] and ketamine [67]. In a similar attempt to connect the latent functional modes to the brain's structure, we recently showed that it was possible to extract the low-dimensional manifold information corresponding to the structural networks using Schroedinger's equation [68]. Similarly, Dan and colleagues developed deep-learning methods to link structural and functional brain dynamics [69, 70]. Altogether, these results show that latent space dimensions are directly connected to structural brain networks, and thus ensure the biological interpretability of our latent whole-brain models, although more work is needed to get a more direct mapping between the two. Further developments in this direction shall shed even more light on our model interpretation shortly.

Latent, biologically relevant, whole-brain models We implemented two different models with different levels of readability: the Hopf and the AHP models. This choice was motivated by the fact that we wanted to make sure that our results would not be model artifacts. This was confirmed given that they are consistent regardless of the model chosen at the node level. However, we also showed that using the AHP model leads to a better-fit quality (2B) while offering a straightforward biological interpretation for each of the fitted parameters (see Figs. 2C and 5B, and the last paragraph of section 1.4). In particular, this new model is paving the way for the development of neuron-glia whole-brain models. Indeed, in the past decades, it has become evident that neuron-only approaches were not sufficient to grasp and explain the complexities of the brain as well as to investigate its disorders. For example, recent modeling efforts have shown that astrocyte calcium activity could explain the neurovascular coupling leading to the BOLD response [33]. Initiating a direction of biologically motivated neuron-glia models *at the whole-brain level* is, to the best of our knowledge, still a niche approach. However, regarding the importance of glial

cell regulation of brain activity, it seems fundamental. In this work, we showed that adding basic astrocyte regulatory mechanisms into our model already could increase our model quality and the data's explainability. In the future, it would be interesting to push this approach forward and work on implementing models that account more explicitly for astrocytes, and other types of glial cells, such as oligodendrocytes and microglia[71, 72].

Finally, we point out that the fMRI processing pipeline provided here could also be used to study different datasets or conditions. Moreover, the modular structure of the provided code and latent whole brain models could even allow one to plug in any other model at the node level that would account for the mechanism one is interested in studying, without having to change the rest of the pipeline. We thus propose here a versatile analysis toolkit that could have multiple applications.

The GEC carries the information on the condition Functional connectivity matrices were already known to be able to distinguish between UWS and MCS [13], and later works on effective connectivity confirmed these findings and showed that even better distinction could be made, in particular in the latent space[51]. The effective connectivity framework also allows us to distinguish between tasks such as rest vs movie watching [73]. In this study, we fitted Generative Effective Connectivity matrices, in a latent space, to analyze how they carried information about the severity of the DoC condition (Fig. 3). In particular, using graph theory, we showed that our subjects' GEC presented two cliques of nodes, strongly connected positively but negatively projecting on the other subgroup (Fig. 3B). Although this structure was similar in different conditions, we found significant differences between the three groups of subjects that we had. This analysis is thus a new step towards understanding in more detail the meso-scale dynamics that are underlying consciousness, and how the GEC matrices carry them. This approach can be related to the circuit models of consciousness [10, 12, 11, 74].

Local node information correlates to prognosis and patient's etiology In the last part of this paper, we showed that looking into the local node parameters (node MBBs) revealed novel information about our patients' conditions (Fig. 5). In particular, we found that these node MBBs, allowed us to group the patients into transverse clusters that are agnostic about the severity of the condition but are correlated to the etiology and evolution of the patients' condition, thus paving the way for novel prognosis tools or biomarkers. Moreover, the cluster separability analysis revealed which parameters were leading to belong to each cluster. Given the direct biological interpretation of the AHP model's parameters, this provides a new hypothesis regarding which mechanism is affected in which group of patients.

Finally, this work opens many possibilities for the investigation of potential treatments. Indeed, it would be very interesting to use the fitted models to see how modifying one specific node parameter affects the cluster to which a patient belongs. For example, if, by affecting one specific parameter (i.e., mechanisms or pharmacological pathway), we can move one patient from a least favorable outcome cluster to a more favorable one, this could be very informative about the treatments to explore for this specific patient. This type of analysis could thus lead to a more precision medicine approach to DoC treatments. Indeed, although many studies exist regarding treatment options both through external neuronal stimulations [19] and pharmacological approaches [11], the wide variability of lesions, and underlying causes for the DoC patients remains a barrier to generalization of such treatments.

3 Methods

3.1 Ethics statement

We comply with all relevant EU legislation such as: The Charter of Fundamental Rights of the EU and notably with the Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation). The present project follows the recommendations of the Declaration of Helsinki – Ethical Principles for Medical Research involving Human Subjects.

Human data is systematically anonymized. Written informed consent to participate in the study was obtained directly from healthy control participants and the legal surrogates of the patients.

Each partner providing data ensured that their work complies fully with all applicable local, government and international laws, regulations and guidelines, including those governing health and safety, data protection, and the use of human subjects and good clinical practice (GCP) including Belgium national legislation implementing the Parliament's Directive 2001/20/ECon GCP, the French national legislation implementing the Parliament's Directive 2001/20/EC on GCP, and the Parliament's Directive 2010/63/EU on the protection of animals used for scientific purposes. We are compliant with the ethical and scientific quality standards reported by the GCP guidelines of the European Union (Guideline for GCP E6(R1)) and 'GCP Directive' 2005/28/EC. All research activities comply with the ethical principles of Horizon 2020, French, Spanish and Belgian ethical regulations, European Union and international legislation, including the Charter of Fundamental Rights of the European Union, the European Convention on Human Rights and its Supplementary Protocols.

The datasets used here are the same ones that we used in previous studies and are fully described in [75, 76].

3.2 Participants

Paris The dataset included 77 brain injury patients hospitalized at Paris Pitié-Salpêtrière. The patients' states of consciousness were evaluated using the Coma Recovery Scale-Revised (CRS-R) by trained clinicians. Patients were classified as being in an unresponsive wakefulness state (UWS) if they demonstrated arousal (such as eye-opening) but lacked any signs of awareness, such as voluntary movement. In contrast, those diagnosed with a minimally conscious state (MCS) exhibited behaviours suggestive of awareness, including visual tracking, response to pain, or consistent command-following. Exclusions from the study were made for individuals with errors in T1 image acquisition (n = 5), significant motion artefacts (n = 7), registration errors (n = 4), or large focal brain lesions (n = 4). As a result, the final cohort comprised 33 patients with MCS (11 females, mean age ± SD: 47.25±20.76 years), 24 patients with UWS (10 females, mean age ± SD: 39.25±16.30 years), and 13 healthy controls (7 females, mean age ± SD: 42.54±13.64 years). This research was approved by the local ethics committee, Comité de Protection des Personnes Ile de France 1 (Paris, France), under the protocol code 'Recherche en soins courants' (NEURODOC protocol, n 2013-A01385-40). Informed consent was obtained from the patient's families, and all procedures were conducted under the Declaration of Helsinki and French regulations.

Liege The dataset involved 35 healthy controls (14 females, mean age±SD: 40±14 years) and 48 patients with DoC. Diagnoses were established following at least five Coma CRS-R assessments conducted by trained clinicians. The final diagnosis, confirmed using Positron Emission Tomography (PET), was determined by the highest level of consciousness recorded. Patients diagnosed with a minimally conscious state (MCS) showed relatively preserved metabolism in the fronto-parietal network, whereas those in an unresponsive wakefulness state (UWS) exhibited bilateral hypometabolism in the same network. The cohort ultimately included 33 patients in MCS (9 females, mean age±SD: 45±16 years) and 15 patients in UWS (6 females, mean age±SD: 47±16

years). The Ethics Committee of the Faculty of Medicine at the University of Liege approved the study protocol, and the research was conducted in compliance with the Helsinki Declaration. Written informed consent was obtained from controls and the legal representatives of the patients.

3.3 MRI data acquisition

Paris MRI images were obtained using two distinct acquisition protocols. The first protocol involved 26 patients and 13 healthy controls, with data collected on a 3T General Electric Signa System. Resting-state T2*-weighted images of the entire brain were acquired using a gradient-echo EPI sequence in an axial orientation (200 volumes, 48 slices, slice thickness: 3 mm, TR/TE: 2400 ms/30 ms, voxel size: 3.4375 × 3.4375 × 3.4375 mm, flip angle: 90°, FOV: 220 mm^2). Additionally, an anatomical T1-weighted MPRAGE sequence was acquired during the same session (154 slices, slice thickness: 1.2 mm, TR/TE: 7.112 ms/3.084 ms, voxel size: 1 × 1 × 1 mm, flip angle: 15°). The second protocol involved MRI data acquisition from 51 patients on a 3T Siemens Skyra System. In this case, T2*-weighted whole-brain resting-state images were recorded with a gradient-echo EPI sequence in an axial orientation (180 volumes, 62 slices, slice thickness: 2.5 mm, TR/TE: 2000 ms/30 ms, voxel size: 2 × 2 × 2 mm, flip angle: 90°, FOV: 240 mm^2, multiband factor: 2). An anatomical volume was also acquired during the same session using a T1-weighted MPRAGE sequence (208 slices, slice thickness: 1.2 mm, TR/TE: 1800 ms/2.35 ms, voxel size: 0.85 × 0.85 × 0.85 mm, flip angle: 8°).

Liege MRI images were obtained using a Siemens 3T Trio scanner (Siemens Inc, Munich, Germany). The acquisition protocol included a gradient echo-planar imaging (EPI) sequence with 32 transverse slices and 300 volumes (TR/TE: 2000 ms/30 ms, flip angle: 78°, voxel size: 3 × 3 × 3 mm, FOV: 192 mm). A structural T1-weighted scan was also acquired with 120 transverse slices (TR: 2300 ms, voxel size: 1.0 × 1.0 × 1.2 mm, flip angle: 9°, FOV: 256 mm).

3.4 Data pre-processing

The resting-state fMRI data was pre-processed using FSL (http://fsl.fmrib.ox.ac.uk/fsl), following the methodology outlined in our previous works [76, 75]. Briefly, the analysis utilized MELODIC (Multivariate Exploratory Linear Optimised Decomposition into Independent Components) [77]. The steps involved discarding the initial five volumes, applying motion correction with MCFLIRT[78], performing brain extraction using BET (Brain Extraction Tool)[79], and spatially smoothing the data with a 5mm FWHM Gaussian kernel. Additional processing steps included rigid-body registration, applying a high-pass filter with a cutoff of 100.0 seconds, and conducting single-session independent component analysis (ICA) with automatic dimensionality estimation. Lesion-related artefacts (in patients) and noise components were removed for each subject using FIX (FMRIB's ICA-based X-noisier)[80]. Finally, the images were co-registered, and FSL tools were used to extract time-series data for each subject, which were mapped to MNI space using the Schaefer [55] and Tian [56] parcellations.

3.5 Paris subset used for the cluster analysis

The preprocessing of the functional MRI data was performed using fMRIPrep (version 22.0.2)[81]. For anatomical T1-weighted scans that displayed visible lesions, lesion masks were created and incorporated into the fMRIPrep pipeline. This preprocessing included a 24-parameter head-motion correction [82], removal of the first four volumes, as well as low-pass and high-pass filtering (0.1 - 0.01 Hz). A total of 26 scans were excluded from the fMRI dataset due to preprocessing failures, such as misalignment or segmentation errors, often caused by patient movement or the size of the

lesions. After applying lenient exclusion criteria, scans with a minimum of 180 volumes and a mean framewise displacement of less than 0.55 mm were retained for analysis.

3.6 Latent space exploration

3.6.1 Datasets combination

In all that follows we combined the datasets from Paris and Liege described above in sections 3.2 to 3.4.
To overcome the problem of different time resolutions between the different datasets, we performed a simple time-interpolation step on the subsets with larger TR to obtain time-series that all have the same $TR = 2$s. This was done using the $interp1d$ function from the *scipy.interpolate* package. The decision of combining the different datasets was taken for the sake of robustness. Indeed, any classification algorithm trained on only one dataset would likely pick up some artifacts or features specific to the scanner used. This could result in overfitting to this specific dataset. Mixing different datasets allows us to avoid this pitfall and ensure the generalizability of our methods, regardless of the center and technical setup where the data was recorded. Furthermore, the preprocessing was strictly the same for all recordings from both centers, thus ensuring no differentiation between the datasets at this stage. Moreover, combining the datasets greatly increases the number of patients in each group and thus the statistical power and significance of any group-related analysis. We recall that the group classification (between MCS and UWS) used in the manuscript comes from the clinical diagnosis performed by the expert clinicians in Paris and Liege based on the CRS-R method which is a universal framework for clinical diagnosis developed for that purpose to ensure that diagnosis of DoC patients is consistent between different centers. Finally, the dimensionality reduction also reduces possible differences due to the differences in data recording.

3.6.2 Auto-encoders

We used auto-encoders with three dense (i.e., fully connected) layers of dimension $0.75 \times d_{init}$, $0.5 \times d_{init}$ and $0.25 \times d_{init}$ before the bottleneck layer, where $d_{init} \in \{100, 116, 1000\}$ is the dimension of the natural space data depending on the initial parcellation (Schaeffer 100 and 1000 [55] and Schaeffer 100 with Tian 16 for the subcortical regions [56]). For each dimension of the bottleneck layer $d_{latent} \in [\![2, 25]\!]$, we split the subjects into a 90-10% training-test split, then each data point given as input to the AE is a vector of dimension d_{init} which corresponds to one time-point in the data time series of one patient. The mean squared reconstruction error (MSE) is then calculated as the MSE between the input and output vectors on the test set. For each d_{latent} we performed a k-fold cross-validation using the *Kfold* function of the Python library *Sckit-learn* (*sklearn* v1.2.2) with $k = 10$.

The auto-encoders were implemented using the *Keras* library of *Tensorflow* v2.11 and trained using *cuda* 11.8 for GPU processing. Training hyperparameters are $LEARNING_RATE = 0.0001$, $BATCH_SIZE = 20$, $EPOCHS = 15$.

3.6.3 Principal Component Analysis

The PCA was implemented in Python using native Python functions and *Numpy*.

3.6.4 SVM classifier

We used the SVC method from the *sklearn* Python library v1.2.2 and leave-one-out cross-validation using the KFold method from *sklearn*. We used a degree 3 polynomial kernel with hyperparameters $C = 0.1$ and $\gamma = 7.5$. The algorithm was trained to distinguish the empirical Functional Connectivity matrices (FC) of the subjects. To enhance the classification accuracy, due to the limited number

of subjects available, we previously split the data time-series into $n_{splits} = 4$ and then computed the FC of the split time-series which led us to $N = n_{splits} \times N_{sub} = 556$ FC matrices to train the SVM classifier, and 4 for the test (we do the Leave-one-out split on the subjects to avoid overfitting by training on some FC matrices of a subject and testing on the same subject).

3.7 Models

3.7.1 Hopf Model

The classical Hopf model is composed, in each node of the model, of the normal form of a supercritical Hopf bifurcation (Stuart-Landau oscillator), close to its bifurcation point. Here we have $n_{nodes} = d_{opt} = 15$ nodes in the latent space. The nodes are then connected via a connectivity matrix C, which we describe in section 3.8. The model equations at node i are given by:

$$\begin{aligned} \dot{x} &= (a - x_i^2 - y_i^2)x_i - \omega_i y_i + G \sum_{j=1}^{d_{opt}} C_{ij}\ (x_j - x_i) + \sigma \eta_i(t) \\ \dot{y} &= (a - x_i^2 - y_i^2)y_i + \omega_i x_i + G \sum_{j=1}^{d_{opt}} C_{ij}\ (y_j - y_i) + \sigma \eta_i(t) \end{aligned} \tag{1}$$

where x is the real part of the system and y its imaginary part, G denotes the global coupling of the system, that we keep set to $G = 1$, C is the connectivity matrix, ω_i is the intrinsic frequency at node i (here we take the Spectral Edge Frequency at 50% (SEF50) value of the spectrogram of the signal) and η_i is an additive Gaussian noise term at each node. We recall that $d_{opt} = 15$ is the dimension of the latent space and thus corresponds to the number of nodes in our model. Finally, a, the bifurcation parameter, and σ the noise intensity are the parameters that we will fit for each patient (see section 3.8).
We chose to keep the global coupling parameter set to $G = 1$ because we fitted the entire connectivity matrix C (section 3.8); thus, any changes in the global coupling intensity are integrated in the fitting of the connectivity matrix. From a mathematical point of view, G being a constant, this is the same as rewriting equation (1) as:

$$\begin{aligned} \dot{x} &= (a - x_i^2 - y_i^2)x_i - \omega_i y_i + \sum_{j=1}^{d_{opt}} J_{ij}\ (x_j - x_i) + \sigma \eta_i(t) \\ \dot{y} &= (a - x_i^2 - y_i^2)y_i + \omega_i x_i + \sum_{j=1}^{d_{opt}} J_{ij}\ (y_j - y_i) + \sigma \eta_i(t) \end{aligned} \tag{2}$$

with $J_{ij} = G \times C_{ij}$.

3.7.2 AHP Model

Local AHP model For the AHP model we implemented a whole-brain version of the model introduced and fully described in [58, 37] which will describe the activity of each node of the model. Here, we recall the model's equations for one node. The model has three variables, the mean population voltage h, the short-term synaptic facilitation x, and the short-term synaptic depression y:

$$\begin{aligned} \tau_0 \dot{h} &= -(h - T_0) + Jxy(h - T_0)^+ + \sqrt{\tau_0}\sigma\dot{\omega} \\ \dot{x} &= \frac{X_{eq} - x}{\tau_f} + K(1 - x)(h - T_0)^+ \\ \dot{y} &= \frac{1 - y}{\tau_r} - Lxy(h - T_0)^+, \end{aligned} \tag{3}$$

where $h^+ = max(h, 0)$ is the population mean firing rate [83, 84], J is the average local connectivity parameter [85, 86], i.e., within one local population (or node), the influence of the other nodes will be added below in a similar way (see equation 4). The parameters K and L describe how the firing rate influences the duration and probability of vesicular release respectively. The time scales τ_f (for facilitation) and τ_r (for depression) define the recovery rates of a synapse from the local network activity and X_{eq} is the resting-state value for the facilitation, which can be associated to the presynaptic calcium concentration. $\dot{\omega}$ is an additive Gaussian noise and σ its amplitude. Finally, we accounted for AHP with two features: 1) a varying equilibrium state T_0 representing hyperpolarization at the end of a burst of neuronal activity 2) two timescales for the medium and slow AHP recovery rates.
The resting membrane potential T_0 and the recovery time constant τ_0 of the voltage h are defined piece-wise as follows:

- $\tau_0 = \tau$ and $T_0 = 0$ when $\{y > Y_{AHP}$ and $h \geq H_{AHP}\}$, represents the fluctuations around the resting membrane potential and the bursting dynamics.

- $\tau_0 = \tau_{mAHP}$ and $T_0 = T_{AHP} < 0$ for $\{y < \dfrac{1}{1 + Lx(h - T_0)} (\iff \dot{y} > 0)$ and $y < Y_h\}$ when the hyperpolarizing currents at the end of the burst become dominant and force the voltage to hyperpolarize (medium AHP period).

- $\tau_0 = \tau_{sAHP}$ and $T_0 = 0$ for $\{y < \dfrac{1}{1 + Lx(h - T_0)}$ and $(Y_{AHP} < y$ or $h < H_{AHP})\}$, which represents the slow recovery to resting membrane potential (slow AHP period).

The threshold parameters defining the three phases are $Y_h = 0.5$, $Y_{AHP} = 0.85$ and $H_{AHP} = -7.5$.

Generalization to a whole-brain model To generalize this model to a whole-brain version, we just need to add the long-range connections, i.e., the influence of the other nodes of the model, as follows. For node i we now have:

$$\begin{aligned} \tau_{0,i}\dot{h}_i &= -(h_i - T_0) + \textstyle\sum_{j=1}^{d_opt} J_{ij}x_jy_j(h_j - T_{0,i})^+ + \sqrt{\tau_{0,i}}\sigma\dot{\omega} \\ \dot{x}_i &= \frac{X_{eq} - x_i}{\tau_f} + K(1 - x_i)(h_i - T_{0,i})^+ \\ \dot{y}_i &= \frac{Y_{eq} - y_i}{\tau_r} - Lx_iy_i(h - T_{0,i})^+. \end{aligned} \tag{4}$$

Similarly as for the Hopf model, the intrinsic timescale τ is extracted directly from the data using the Spectral Edge Frequency at 50% (SEF50) from the spectrogram of the signal. Similarly as for the Hopf model, we directly fit a connectivity matrix $J_{ij} = G \times C_{ij}$, integrating the global coupling into the fitted connectivity matrix. Note that we have replaced the resting-state value of the depression that was left fixed at 1 in [58, 37, 87, 84] by a parameter Y_{eq} that we will fit to the data to be able to grasp potentially depressed synapses in some patients. The fixed model parameters are taken from [58] and are given in Table 2.

	Parameters	Values
K	Facilitation rate	0.037Hz [84, 58]
L	Depression rate	0.028Hz [84, 58]
τ_r	Facilitation timescale	2.9s [58]
τ_f	Depression timescale	0.9s [58]
T	voltage resting-state	0

Table 2: Fixed AHP model parameters

3.7.3 Balloon-Windkessel model

To simulate the BOLD signal, we add a Balloon-Windkessel transformation as described in [59, 60]. We used the implementation from github.com/dagush/WholeBrain/

3.7.4 Models implementation and numerical integration

All models are implemented in Python, with Numba (v0.56.4), a just-in-time compiler, to speed up the computations. The numerical integration for the Hopf and AHP model is done with the Runge-Kutta 4 scheme with a time step of $dt_{sim} = 0.05$s.

3.8 Fitting algorithm

The fitting algorithm is defined in two recurrent steps (see Fig. 2A). Both steps are based on maximizing the structural similarity index measure (SSIM) between empirical and simulated connectivity matrices

SSIM The SSIM index is a measure of similarity between two images (or matrices) of the same size [61]. It is computed by comparing common-size windows of corresponding zones of both matrices and then aggregating the results. The formula used for the comparison of two matrices A and B is:

$$SSIM(A,B) = \frac{(2\mu_A\mu_B + c_1)\,(2\sigma_{AB} + c_2)}{\left(\mu_A^2 + \mu_B^2 + c_1\right)\left(\sigma_A^2 + \sigma_B^2 + c_2\right)}, \tag{5}$$

where:
μ_X is the pixel empirical mean of matrix $X \in A, B$ and σ_X its empirical variance;
σ_{AB} the covariance of the two matrices;
c_1 and c_2 are two variables introduced to stabilize the division with a weak denominator, we used the classical default values:
$c_1 = (0.01L)^2$ and $c_2 = (0.03L)^2$ where L is the dynamic range of pixel values in A and B, which in our case is $L = 1.0$.

Node parameters fitting First we fit the node parameters, that is $\Theta = a, \sigma$ for the Hopf model and $\Theta = \{\sigma, X_{eq}, Y_{eq}, T_h, 1/\tau_{mAHP}, 1/\tau_{sAHP}\}$ for the AHP model. The goal of this step is to find the choice of node parameters Θ that will maximize the SSIM between empirical and simulated functional connectivity matrices for each patient. To do so, we keep the connectivity matrix fixed and vary the node parameters Θ exploring the parameter space.
We explore randomly the parameter space in the manner of a multi-scale grid search, i.e., we first explore a larger range of parameters and at each iteration, we refine the search around the best-found value. In practice, at the first iteration $n_{iter} = 0$, the connectivity matrix C is defined as the empirical functional connectivity matrix, FC_{emp}, taken from the data, and for the following iterations, we take $C_{iter} = GEC_{iter-1}$, where $iter$ is the iteration step of the fitting procedure. Here, we stopped at $iter_{max} = 4$ for the Hopf model and $iter_{max} = 6$ for the AHP model, when it was clear that the SSIM stopped increasing with fit iterations.
At each iteration $iter$, we explored $N_{max} = (4 - iter) \times 1000$ (for $iter \leq 3$ and then $N_{max} = 1000$ for the following iterations) values for the node parameters Θ. Since we have a stochastic model we need to run many simulations of the model at each iteration. Thus, for each value of Θ we ran $N_{sim} = (iter \times 10) + 15$ (for $iter \leq 3$ and then $N_{sim} = 50$ for the following iterations) simulations of duration $dur = 300$s (same as our recordings) and then averaged them before taking the correlation matrix to obtain the simulated FC FC_{sim}. The reason why the number of simulations ran depends on the iteration step $iter$ is to mimic the approach of multi-level grid searches where we first explore a wider range of parameters (i.e., a bigger N_{max}) with less precision (i.e., a smaller N_{sim}) and then

progressively explore fewer parameters (decreasing N_{max}) with more precision (increasing N_{sim}) while keeping reasonable computation times.
The explored parameter ranges are summarized in Table 3.

Generative Effective Connectivity fitting Briefly, the generative effective connectivity matrix (GEC) is a matrix that allows for the reproduction of both the empirical correlation matrix of the data and a forward time-shifted correlation matrix. This is done to grasp the asymmetries in the effective connectivity between the nodes and leads to an asymmetric matrix which can reproduce non-equilibrium dynamics (see, for example, [51, 88] for more details about GEC matrices). In practice, we fitted the GEC by keeping the node parameters from the previous step Θ fixed and adjusting the connectivity matrix to maximize the SSIM between the empirical FC^{emp} vs simulated FC^{sim} matrices, as well as between the shifted matrices $FC^{emp-tau}$ and $FC^{sim-tau}$. The shifted matrices were obtained by shifting in time each of the time series of a small time step τ and taking its correlation with the unshifted other dimensions. This process is done to grasp the asymmetries in the effective connectivity between the nodes. The GEC was fitted iteratively using a pseudo gradient descent algorithm according to equation (6) below, which weighs both the fitting of the direct and shifted FC matrices, as also explained in [51]:

$$GEC_{ij} = GEC_{ij} + \epsilon(FC^{emp} - FC^{sim} + FC^{emp-tau} - FC^{sim-tau}). \tag{6}$$

The learning rate is $\epsilon = 0.01$ and the shift in time $\tau = 3 \times TR = 6$s [51]. Similarly, as in the node fitting step, we need to run several simulations at each step to account for the stochasticity, here we have $N_{sim} = 5$. For each iteration, this is repeated until the SSIM is larger than 0.95 or for $N_{max} = 1000$ iterations.

	All	Hopf	AHP model				
	σ	a	X_{eq}	Y_{eq}	T_h	τ_{mAHP}	τ_{sAHP}
Initial value	5	-0.02	0.088	1	-30	0.15s	5s
Range	$[0.5, 7.5]$	$[-0.5, 0.5]$	$[0.01, 0.21]$	$[0.1, 1]$	$[-50, -10]$	$[0.05, 0.5]$s	$[0.75, 15]$s

Table 3: Authorized parameter ranges

3.9 GEC and Graph measures

All the graphs are constructed with the $from_numpy_matrix$ function of the package *networkx* (v2.8.4). We then split all the graphs between two subgraphs, one containing all the positive edges, and another one containing only the negative edges. For the interclique analysis (Fig. 3C.) we use the $find_cliques$ function to get the maximum clique in which one particular node is involved and take all the other nodes (the ones not contained in the maximal clique) as the other clique, and then computed the average and standard deviation of the connectivity between these two groups of nodes using the *edge_boundary* function, this is repeated for all nodes. The other graph measures (Fig. 3D.) are directly obtained with built-in *networkx* functions for node centrality and clustering coefficients.
The GEC quantifications (Fig. 3E.) are computed as follow:

$$sparsity = \frac{n_{small}}{d_{opt}^2}, \tag{7}$$

where n_{small} is the number of values in GEC that are smaller (in absolute value) than 0.015.

$$asymmetry = \frac{|GEC - GEC^T|}{2} \tag{8}$$

The mean connectivity and standard deviation are the mean value of the GEC and its standard deviation (computed with the *numpy* built-in functions).

3.10 Decoding of the GEC cliques and latent dimensions

To understand the connection between the latent dimensions and the natural brain structures, we created surrogate signals in the natural space corresponding to each latent dimension, as follows:

1. We projected all the control subjects' data into the latent space, using the same trained AE as for the rest of the study.

2. From each encoded subject, we created $d_opt = 15$ surrogate signals, by keeping the encoded data on one of the latent dimensions $LD \in [|0, 14|]$ and setting the other ones to zero.

3. We decoded the surrogate signals (with the decoder part of the same trained AE) to project them back into the natural space. At this stage, we obtained 15 surrogate fMRI signal (with 100 time-series − the dimension of the natural space) per control subject.

4. We averaged the decoded fMRI signals over all control subjects. This was done to strengthen the results' robustness, after we verified that, as expected, all decoded signals per latent dimension had a very similar structure, regardless of the dataset (Paris or Liege).

Then we compared the obtained signals for each latent dimension to the 7 resting state networks (RSNs) from [54]. To do so, we computed a FC matrix for each RSN, in the natural space, by keeping the control subjects' data only in the nodes corresponding to one of the RSNs and putting very small amplitude white noise on the other dimensions (generated with $10^{-3} \times numpy.random.rand$). The small noise on the remaining dimensions was added to to avoid the division by zero in the computation of the correlation coefficient for signals that would otherwise have a variance equal to zero.
Finally, we computed the SSIM between the decoded FC matrices for each latent dimension ($LD \in [|0, 14|]$) and the FC matrices obtained for each of the 7 RSNs.
We applied the same methodology for the two cliques of the GEC observed in Fig. 3B. This means that instead of keeping the encoded signal only one latent dimension at a time, we kept the signal on all the nodes of one clique and then of the other one. The rest of the analysis process wa kept the same.

3.11 Local node MBBs clustering

3.11.1 Outcome categorization

To analyze the link between the MBB clusters and the patients metadata, we defined five outcome categories for the patients based on the successive CRS-R diagnosis available. The final outcome available in our dataset is the CRS-R diagnosis after two years unless the patient died or regained consciousness before that.

1. Dead;

2. Decreasing: when the evolution of the CRS-R diagnosis went towards less conscious over successive checks but did not lead to death;

3. Stable: when the CRS-R diagnosis did not change over time;

4. Increasing: when the diagnosis ameliorated without leading to consciousness;

5. Conscious or Emerged.

We refer to outcomes 1 to 3 as *negative* outcomes, and outcomes 4 and 5 as *positive* or *favorable* outcomes.

3.11.2 UMAPS and HDBSCAN clustering

To obtain the clusters we used Uniform Manifold Approximation and Projection for Dimension Reduction (UMAP)[63] on the 6-dimensional (resp. 2-d) vectors of fitted node parameters (node MBBs) from the AHP (resp. Hopf) model, using the UMAP Python package (*umap-learn* v0.5.5), we used the parameters $n_neighbors = 5$ and $min_dist = 0.1$, to project in 2D, the rest are the default parameters.
The clustering was then done on the UMAP projections using the *HDBScan* clustering algorithm[64] implemented for Python in the *HDBSCAN* package (*hdbscan* v0.8.33). We used parameters $min_samples = 3$ and $min_cluster_size = 3$, and the rest take default values.

3.12 Code availability

All codes are implemented in Python and are available here: https://github.com/louzonca/DoC_models.

4 Authors contribution

Conceptualization: LZ, GD, JDS – data collection: JA, OG, JDS, SL – data pre-processing: AE, DM, LZ – data analysis: LZ – software implementation: LZ, GP – results analysis and discussion: all authors – writing - original draft: LZ – writing - final draft: all authors

5 Acknowledgmenents

LZ was supported by the FLAG-ERA JTC2021 project ModelDXConsciousness (Human Brain Project Partnering Project) Grant PCI2021-122019-2A funded by MICIU/AEI /10.13039/501100011033 and by the European Union NextGenerationEU/PRTR. AE was supported by the project eBRAIN-Health—Actionable Multilevel Health Data (id 101058516), funded by EU Horizon Europe and by the Grant PID2022-136216NBI00, funded by MICIU/AEI/10.13039/501100011033, and "ERDF A way of making Europe", ERDF, EU. YSP was supported by the project NEurological MEchanismS of Injury, and the project Sleep-like cellular dynamics (NEMESIS) (ref. 101071900) funded by the EU ERC Synergy Horizon Europe. GD was supported by the project NEurological MEchanismS of Injury, and the project Sleep-like cellular dynamics (NEMESIS) (ref. 101071900) funded by the EU ERC Synergy Horizon Europe and and by the Grant PID2022-136216NB-I00, funded by MICIU/AEI/10.13039/501100011033, and "ERDF A way of making Europe", ERDF, EU.
This research was partially funded by Grant PID2021-122136OB-C22 funded by MICIU/AEI/10.13039/501100011033 and by ERDF A way of making Europe for GP.
DM received individual funding from Ecole Doctorale Frontières de l'Innovation en Recherche et Education–Programme Bettencourt. DM and JDS were supported by funding from the EU ERAPerMed Joint Translational Call for Proposals for "Personalised Medicine: Multidisciplinary research towards implementation" (ERA PerMed JTC2019). JDS was also supported by the FLAG-ERA JTC2021 project ModelDXConsciousness (Human Brain Project Partnering Project).
The study was further supported by the University and University Hospital of Liège, the BIAL Foundation, the Belgian National Funds for Scientific Research (FRS-FNRS), the FNRS PDR project (T.0134.21), the FLAG-ERA JTC2021 project ModelDXConsciousness (Human Brain Project Partnering Project) and FLAG-ERA JTC 2023 - HBP - Basic and Applied Research, project BrainAct, JTC the fund Generet, the King Baudouin Foundation, the BIAL Foundation, the Mind Science Foundation, the European Commission, the Fondation Leon Fredericq, the Mind-Care foundation, the National Natural Science Foundation of China (Joint Research Project 81471100) and the European Foundation of Biomedical Research FERB Onlus, the Horizon 2020 MSCA – Research and

Innovation Staff Exchange DoC-Box project (HORIZON-MSCA-2022-SE-01-01; 101131344). OG is Research Associate and SL Research Director at FRS-FNRS. JA is postdoctoral fellow funded (1265522N) by the Fund for Scientific Research-Flanders (FWO).

References

[1] Abhidharma chitta dhammas https://en.wikipedia.org/wiki/abhidharma.

[2] Javier Alejandro Galadí. The mind-body problem: An overview of proposed solutions. *The Theory of Mind Under Scrutiny: Psychopathology, Neuroscience, Philosophy of Mind and Artificial Intelligence*, pages 435–467, 2024.

[3] Edwin G Boring. *The physical dimensions of consciousness.* The Century Co., 1933.

[4] Steven Laureys, Adrian M Owen, and Nicholas D Schiff. Brain function in coma, vegetative state, and related disorders. *The Lancet Neurology*, 3(9):537–546, 2004.

[5] Bryan Jennett and Fred Plum. Persistent vegetative state after brain damage: a syndrome in search of a name. *The Lancet*, 299(7753):734–737, 1972.

[6] Joseph T Giacino and Kathleen Kalmar. Diagnostic and prognostic guidelines for the vegwive and minimally conscious states. *Neuropsychological Rehabilitation*, 15(3-4):166–174, 2005.

[7] Marie-Aurélie Bruno, Audrey Vanhaudenhuyse, Aurore Thibaut, Gustave Moonen, and Steven Laureys. From unresponsive wakefulness to minimally conscious plus and functional locked-in syndromes: recent advances in our understanding of disorders of consciousness. *Journal of neurology*, 258:1373–1384, 2011.

[8] Nicholas D Schiff. Cognitive motor dissociation following severe brain injuries. *JAMA neurology*, 72(12):1413–1415, 2015.

[9] Steven Laureys and Nicholas D Schiff. Coma and consciousness: paradigms (re) framed by neuroimaging. *Neuroimage*, 61(2):478–491, 2012.

[10] Nicholas D Schiff. Mesocircuit mechanisms in the diagnosis and treatment of disorders of consciousness. *La Presse Médicale*, 52(2):104161, 2023.

[11] Megan E Barra, Ken Solt, Xin Yu, and Brian L Edlow. Restoring consciousness with pharmacologic therapy: Mechanisms, targets, and future directions. *Neurotherapeutics*, page e00374, 2024.

[12] Min Zhuo. Cortical synaptic basis of consciousness. *European Journal of Neuroscience*, 59(5):796–806, 2024.

[13] Athena Demertzi, Georgios Antonopoulos, Lizette Heine, Henning U Voss, Julia Sophia Crone, Carlo de Los Angeles, Mohamed Ali Bahri, Carol Di Perri, Audrey Vanhaudenhuyse, Vanessa Charland-Verville, et al. Intrinsic functional connectivity differentiates minimally conscious from unresponsive patients. *Brain*, 138(9):2619–2631, 2015.

[14] Johan Stender, Olivia Gosseries, Marie-Aurélie Bruno, Vanessa Charland-Verville, Audrey Vanhaudenhuyse, Athena Demertzi, Camille Chatelle, Marie Thonnard, Aurore Thibaut, Lizette Heine, et al. Diagnostic precision of pet imaging and functional mri in disorders of consciousness: a clinical validation study. *The Lancet*, 384(9942):514–522, 2014.

[15] Yelena G Bodien, Judith Allanson, Paolo Cardone, Arthur Bonhomme, Jerina Carmona, Camille Chatelle, Srivas Chennu, Mary Conte, Stanislas Dehaene, Paola Finoia, et al. Cognitive motor dissociation in disorders of consciousness. *New England Journal of Medicine*, 391(7):598–608, 2024.

[16] Jennifer Egbebike, Qi Shen, Kevin Doyle, Caroline A Der-Nigoghossian, Lucy Panicker, Ian Jerome Gonzales, Lauren Grobois, Jerina C Carmona, Athina Vrosgou, Arshneil Kaur, et al. Cognitive-motor dissociation and time to functional recovery in patients with acute brain injury in the usa: a prospective observational cohort study. *The Lancet Neurology*, 21(8):704–713, 2022.

[17] Aurore Thibaut, Nicholas Schiff, Joseph Giacino, Steven Laureys, and Olivia Gosseries. Therapeutic interventions in patients with prolonged disorders of consciousness. *The Lancet Neurology*, 18(6):600–614, 2019.

[18] Mohammed Saif, Shatha A Sharbatti, Anas Nemmar, Sharanya S Kumar, Krishna Prasad, Asma M Khan, Iman Khadar, Sharmila Banu, Shatha Sharbatti, and Iman Mohamoud. Outcomes of neurorehabilitation among patients with prolonged disorders of consciousness. *Cureus*, 15(5), 2023.

[19] Amber R Hopkins, Marie M Vitello, Aurore Thibaut, and Martin M Monti. Emerging treatment for patients with disorders of consciousness: The field of neuromodulation. In *Coma and Disorders of Consciousness*, pages 147–208. Springer, 2024.

[20] Gustavo Deco, Viktor K Jirsa, and Anthony R McIntosh. Resting brains never rest: computational insights into potential cognitive architectures. *Trends in neurosciences*, 36(5):268–274, 2013.

[21] Gustavo Deco, Adrián Ponce-Alvarez, Patric Hagmann, Gian Luca Romani, Dante Mantini, and Maurizio Corbetta. How local excitation–inhibition ratio impacts the whole brain dynamics. *Journal of Neuroscience*, 34(23):7886–7898, 2014.

[22] Gustavo Deco, Anthony R McIntosh, Kelly Shen, R Matthew Hutchison, Ravi S Menon, Stefan Everling, Patric Hagmann, and Viktor K Jirsa. Identification of optimal structural connectivity using functional connectivity and neural modeling. *Journal of Neuroscience*, 34(23):7910–7916, 2014.

[23] Gustavo Deco and Viktor K Jirsa. Ongoing cortical activity at rest: criticality, multistability, and ghost attractors. *Journal of Neuroscience*, 32(10):3366–3375, 2012.

[24] Rajanikant Panda, Aurore Thibaut, Ane Lopez-Gonzalez, Anira Escrichs, Mohamed Ali Bahri, Arjan Hillebrand, Gustavo Deco, Steven Laureys, Olivia Gosseries, Jitka Annen, et al. Disruption in structural–functional network repertoire and time-resolved subcortical fronto-temporoparietal connectivity in disorders of consciousness. *Elife*, 11:e77462, 2022.

[25] Joana Cabral, Henry Luckhoo, Mark Woolrich, Morten Joensson, Hamid Mohseni, Adam Baker, Morten L Kringelbach, and Gustavo Deco. Exploring mechanisms of spontaneous functional connectivity in meg: how delayed network interactions lead to structured amplitude envelopes of band-pass filtered oscillations. *Neuroimage*, 90:423–435, 2014.

[26] Ben A Barres. The mystery and magic of glia: a perspective on their roles in health and disease. *Neuron*, 60(3):430–440, 2008.

[27] Nicola J Allen and Ben A Barres. Glia—more than just brain glue. *Nature*, 457(7230):675–677, 2009.

[28] Gertrudis Perea, Marta Navarrete, and Alfonso Araque. Tripartite synapses: astrocytes process and control synaptic information. *Trends in neurosciences*, 32(8):421–431, 2009.

[29] Glenn Dallérac, Oana Chever, and Nathalie Rouach. How do astrocytes shape synaptic transmission? insights from electrophysiology. *Frontiers in cellular neuroscience*, 7:159, 2013.

[30] Oana Chever, Elena Dossi, Ulrike Pannasch, Mickael Derangeon, and Nathalie Rouach. Astroglial networks promote neuronal coordination. *Science signaling*, 9(410), January 2016.

[31] Anke Wallraff, Rüdiger Köhling, Uwe Heinemann, Martin Theis, Klaus Willecke, and Christian Steinhäuser. The impact of astrocytic gap junctional coupling on potassium buffering in the hippocampus. *Journal of Neuroscience*, 26(20):5438–5447, 2006.

[32] Ulrike Pannasch, Mickael Derangeon, Oana Chever, and Nathalie Rouach. Astroglial gap junctions shape neuronal network activity. *Communicative & integrative biology*, 5(3):248–254, 2012.

[33] Federico Tesler, Marja-Leena Linne, and Alain Destexhe. Modeling the relationship between neuronal activity and the bold signal: contributions from astrocyte calcium dynamics. *Scientific Reports*, 13(1):6451, 2023.

[34] Philip G Haydon and Giorgio Carmignoto. Astrocyte control of synaptic transmission and neurovascular coupling. *Physiological reviews*, 86(3):1009–1031, 2006.

[35] Jessica A Filosa and Jennifer A Iddings. Astrocyte regulation of cerebral vascular tone. *American Journal of Physiology-Heart and Circulatory Physiology*, 305(5):H609–H619, 2013.

[36] Costantino Iadecola. The neurovascular unit coming of age: a journey through neurovascular coupling in health and disease. *Neuron*, 96(1):17–42, 2017.

[37] Elena Dossi, Lou Zonca, Helena Pivonkova, Giampaolo Milior, Julien Moulard, Lydia Vargova, Oana Chever, David Holcman, and Nathalie Rouach. Astroglial gap junctions strengthen hippocampal network activity by sustaining afterhyperpolarization via kcnq channels. *Cell Reports*, 43(5), 2024.

[38] Jeremie Sibille, Khanh Dao Duc, David Holcman, and Nathalie Rouach. The neuroglial potassium cycle during neurotransmission: Role of kir4.1 channels. *PLoS Computational Biology*, 11(3), March 2015.

[39] Nathalie Rouach, K Dao Duc, Jeremie Sibille, and David Holcman. Dynamics of ion fluxes between neurons, astrocytes and the extracellular space during neurotransmission. *Opera Medica et Physiologica*, 4(1), 2018.

[40] Marc R Freeman and David H Rowitch. Evolving concepts of gliogenesis: a look way back and ahead to the next 25 years. *Neuron*, 80(3):613–623, 2013.

[41] Alexei Verkhratsky and Arthur Butt. *Glial physiology and pathophysiology*. John Wiley & Sons, 2013.

[42] Ashley M Ingiosi, Christopher R Hayworth, Daniel O Harvey, Kristan G Singletary, Michael J Rempe, Jonathan P Wisor, and Marcos G Frank. A role for astroglial calcium in mammalian sleep and sleep regulation. *Current Biology*, 30(22):4373–4383, 2020.

[43] Laura Bojarskaite, Daniel M Bjørnstad, Klas H Pettersen, Céline Cunen, Gudmund Horn Hermansen, Knut Sindre Åbjørsbråten, Anna R Chambers, Rolf Sprengel, Koen Vervaeke, Wannan Tang, et al. Astrocytic ca^{2+} signaling is reduced during sleep and is involved in the regulation of slow wave sleep. *Nature communications*, 11(1):3240, 2020.

[44] In-Sun Choi, Jae-Hong Kim, Ji-Young Jeong, Maan-Gee Lee, Kyoungho Suk, and Il-Sung Jang. Astrocyte-derived adenosine excites sleep-promoting neurons in the ventrolateral preoptic nucleus: Astrocyte-neuron interactions in the regulation of sleep. *Glia*, 70(10):1864–1885, 2022.

[45] Ashley M Ingiosi and Marcos G Frank. Goodnight, astrocyte: waking up to astroglial mechanisms in sleep. *The FEBS journal*, 290(10):2553–2564, 2023.

[46] Peiran Gao and Surya Ganguli. On simplicity and complexity in the brave new world of large-scale neuroscience. *Current opinion in neurobiology*, 32:148–155, 2015.

[47] Mehrdad Jazayeri and Arash Afraz. Navigating the neural space in search of the neural code. *Neuron*, 93(5):1003–1014, 2017.

[48] James M Shine, Michael Breakspear, Peter T Bell, Kaylena A Ehgoetz Martens, Richard Shine, Oluwasanmi Koyejo, Olaf Sporns, and Russell A Poldrack. Human cognition involves the dynamic integration of neural activity and neuromodulatory systems. *Nature neuroscience*, 22(2):289–296, 2019.

[49] Joan Rué-Queralt, Angus Stevner, Enzo Tagliazucchi, Helmut Laufs, Morten L Kringelbach, Gustavo Deco, and Selen Atasoy. Decoding brain states on the intrinsic manifold of human brain dynamics across wakefulness and sleep. *Communications Biology*, 4(1):854, 2021.

[50] Gustavo Deco, Diego Vidaurre, and Morten L Kringelbach. Revisiting the global workspace orchestrating the hierarchical organization of the human brain. *Nature human behaviour*, 5(4):497–511, 2021.

[51] Yonatan Sanz Perl, Sebastian Geli, Eider Pérez-Ordoyo, Lou Zonca, Sebastian Idesis, Jakub Vohryzek, Viktor K. Jirsa, Morten L. Kringelbach, Enzo Tagliazucchi, and Gustavo Deco. Modelling low-dimensional interacting brain networks reveals organising principle in human cognition. *Network Neuroscience*, pages 1–21, 01 2025.

[52] Sebastian Idesis, Michele Allegra, Jakub Vohryzek, Yonatan Sanz Perl, Joshua Faskowitz, Olaf Sporns, Maurizio Corbetta, and Gustavo Deco. A low dimensional embedding of brain dynamics enhances diagnostic accuracy and behavioral prediction in stroke. *Scientific Reports*, 13(1):15698, 2023.

[53] Yonatan Sanz Perl, Carla Pallavicini, Juan Piccinini, Athena Demertzi, Vincent Bonhomme, Charlotte Martial, Rajanikant Panda, Naji Alnagger, Jitka Annen, Olivia Gosseries, et al. Low-dimensional organization of global brain states of reduced consciousness. *Cell Reports*, 42(5), 2023.

[54] BT Thomas Yeo, Fenna M Krienen, Jorge Sepulcre, Mert R Sabuncu, Danial Lashkari, Marisa Hollinshead, Joshua L Roffman, Jordan W Smoller, Lilla Zöllei, Jonathan R Polimeni, et al. The organization of the human cerebral cortex estimated by intrinsic functional connectivity. *Journal of neurophysiology*, 2011.

[55] Alexander Schaefer, Ru Kong, Evan M Gordon, Timothy O Laumann, Xi-Nian Zuo, Avram J Holmes, Simon B Eickhoff, and BT Thomas Yeo. Local-global parcellation of the human cerebral cortex from intrinsic functional connectivity mri. *Cerebral cortex*, 28(9):3095–3114, 2018.

[56] Ye Tian, Daniel S Margulies, Michael Breakspear, and Andrew Zalesky. Topographic organization of the human subcortex unveiled with functional connectivity gradients. *Nature neuroscience*, 23(11):1421–1432, 2020.

[57] Gustavo Deco, Morten L Kringelbach, Viktor K Jirsa, and Petra Ritter. The dynamics of resting fluctuations in the brain: metastability and its dynamical cortical core. *Scientific reports*, 7(1):3095, 2017.

[58] Lou Zonca and David Holcman. Modeling bursting in neuronal networks using facilitation-depression and afterhyperpolarization. *Communications in Nonlinear Science and Numerical Simulation*, 94:105555, 2021.

[59] Klaas Enno Stephan, Nikolaus Weiskopf, Peter M. Drysdale, Peter A. Robinson, and Karl J. Friston. Comparing hemodynamic models with dcm. *NeuroImage*, 38(3):387–401, 2007.

[60] Klaas Enno Stephan, Lars Kasper, Lee M. Harrison, Jean Daunizeau, Hanneke E.M. den Ouden, Michael Breakspear, and Karl J. Friston. Nonlinear dynamic causal models for fmri. *NeuroImage*, 42(2):649–662, 2008.

[61] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.

[62] Ignacio Perez Ipiña, Patricio Donnelly Kehoe, Morten Kringelbach, Helmut Laufs, Agustín Ibañez, Gustavo Deco, Yonatan Sanz Perl, and Enzo Tagliazucchi. Modeling regional changes in dynamic stability during sleep and wakefulness. *NeuroImage*, 215:116833, 2020.

[63] Leland McInnes, John Healy, and James Melville. Umap: Uniform manifold approximation and projection for dimension reduction. *arXiv preprint arXiv:1802.03426*, 2018.

[64] Myeong-Hun Jeong, Yaping Cai, Clair J Sullivan, and Shaowen Wang. Data depth based clustering analysis. In *Proceedings of the 24th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, pages 1–10, 2016.

[65] Selen Atasoy, Gustavo Deco, Morten L Kringelbach, and Joel Pearson. Harmonic brain modes: a unifying framework for linking space and time in brain dynamics. *The Neuroscientist*, 24(3):277–293, 2018.

[66] Jakub Vohryzek, Joana Cabral, Christopher Timmermann, Selen Atasoy, Leor Roseman, David J Nutt, Robin L Carhart-Harris, Gustavo Deco, and Morten L Kringelbach. The flattening of spacetime hierarchy of the n, n-dimethyltryptamine brain state is characterized by harmonic decomposition of spacetime (hades) framework. *National Science Review*, 11(5):nwae124, 2024.

[67] Milan Van Maldegem, Jakub Vohryzek, Selen Atasoy, Naji Alnagger, Paolo Cardone, Vincent Bonhomme, Audrey Vanhaudenhuyse, Athena Demertzi, Oceane Jaquet, Mohamed Ali Bahri, et al. Ketamine-induced unresponsiveness shows a harmonic shift from global to localised functional organisation. *bioRxiv*, pages 2024–06, 2024.

[68] Gustavo Deco, Yonatan Sanz Perl, and Morten L Kringelbach. Complex harmonics reveal low-dimensional manifolds of critical brain dynamics. *bioRxiv*, pages 2024–06, 2024.

[69] Tingting Dan, Zhuobin Huang, Hongmin Cai, Robert G Lyday, Paul J Laurienti, and Guorong Wu. Uncovering shape signatures of resting-state functional connectivity by geometric deep learning on riemannian manifold. *Human Brain Mapping*, 43(13):3970–3986, 2022.

[70] Tingting Dan, Ziquan Wei, Won Hwa Kim, and Guorong Wu. Exploring the enigma of neural dynamics through a scattering-transform mixer landscape for riemannian manifold. *arXiv preprint arXiv:2405.16357*, 2024.

[71] Alireza Poshtkohi, John Wade, Liam McDaid, Junxiu Liu, Mark Dallas, and Angela Bithell. Mathematical modelling of human p2x-mediated plasma membrane electrophysiology and calcium dynamics in microglia. *PLOS Computational Biology*, 17(11):e1009520, 2021.

[72] Leah E Vaughan, Prerna R Ranganathan, Raj G Kumar, Amy K Wagner, and Jonathan E Rubin. A mathematical model of neuroinflammation in severe clinical traumatic brain injury. *Journal of neuroinflammation*, 15:1–19, 2018.

[73] Matthieu Gilson, Gustavo Deco, Karl J Friston, Patric Hagmann, Dante Mantini, Viviana Betti, Gian Luca Romani, and Maurizio Corbetta. Effective connectivity inferred from fmri transition dynamics during movie viewing points to a balanced reconfiguration of cortical interactions. *Neuroimage*, 180:534–546, 2018.

[74] Prejaas Tewarie, Romesh Abeysuriya, Rajanikant Panda, Pablo Nunez, Marie M Vitello, Glenn van der Lande, Olivia Gosseries, Aurore Thibaut, Steven Laureys, Gustavo Deco, et al. Individual trajectories for recovery of neocortical activity in disorders of consciousness. *bioRxiv*, pages 2024–03, 2024.

[75] Ane López-González, Rajanikant Panda, Adrián Ponce-Alvarez, Gorka Zamora-López, Anira Escrichs, Charlotte Martial, Aurore Thibaut, Olivia Gosseries, Morten L Kringelbach, Jitka Annen, et al. Loss of consciousness reduces the stability of brain hubs and the heterogeneity of brain dynamics. *Communications biology*, 4(1):1037, 2021.

[76] Anira Escrichs, Yonatan Sanz Perl, Carme Uribe, Estela Camara, Basak Türker, Nadya Pyatigorskaya, Ane López-González, Carla Pallavicini, Rajanikant Panda, Jitka Annen, et al. Unifying turbulent dynamics framework distinguishes different brain states. *Communications Biology*, 5(1):638, 2022.

[77] Christian F Beckmann and Stephen M Smith. Probabilistic independent component analysis for functional magnetic resonance imaging. *IEEE transactions on medical imaging*, 23(2):137–152, 2004.

[78] Mark Jenkinson, Peter Bannister, Michael Brady, and Stephen Smith. Improved optimization for the robust and accurate linear registration and motion correction of brain images. *Neuroimage*, 17(2):825–841, 2002.

[79] Stephen M Smith. Fast robust automated brain extraction. *Human brain mapping*, 17(3):143–155, 2002.

[80] Ludovica Griffanti, Gholamreza Salimi-Khorshidi, Christian F Beckmann, Edward J Auerbach, Gwenaëlle Douaud, Claire E Sexton, Enikő Zsoldos, Klaus P Ebmeier, Nicola Filippini, Clare E Mackay, et al. Ica-based artefact removal and accelerated fmri acquisition for improved resting state network imaging. *Neuroimage*, 95:232–247, 2014.

[81] Oscar Esteban, Christopher J Markiewicz, Ross W Blair, Craig A Moodie, A Ilkay Isik, Asier Erramuzpe, James D Kent, Mathias Goncalves, Elizabeth DuPre, Madeleine Snyder, et al. fmriprep: a robust preprocessing pipeline for functional mri. *Nature methods*, 16(1):111–116, 2019.

[82] Marina Weiler, Raphael F Casseb, Brunno M de Campos, Julia S Crone, Evan S Lutkenhoff, Paul M Vespa, Martin M Monti, and EpiBioS4Rx Study Group. Evaluating denoising strategies in resting-state functional magnetic resonance in traumatic brain injury (epibios4rx). *Human brain mapping*, 43(15):4640–4649, 2022.

[83] David Holcman and Misha Tsodyks. The emergence of up and down states in cortical networks. *PLoS Computational Biology*, 2(3):174–181, March 2006.

[84] K Dao Duc, Chun-Yao Lee, Pierre Parutto, Dror Cohen, Menahem Segal, Nathalie Rouach, and David Holcman. Bursting reverberation as a multiscale neuronal network process driven by synaptic depression-facilitation. *PloS One*, 10(5):e0124694, 2015.

[85] Misha V. Tsodyks and Henry Markram. The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability. *Proc. Natl. Acad. Sci. USA*, 94:719–723, January 1997.

[86] Evgeniy Bart, Shaowen Bao, and David Holcman. Modeling the spontaneous activity of the auditory cortex. *Journal of computational neuroscience*, 19(3):357–378, 2005.

[87] Lou Zonca and David Holcman. Emergence and fragmentation of the alpha-band driven by neuronal network dynamics. *PLoS Computational Biology*, 17(12):e1009639, 2021.

[88] Sebastian Idesis, Michele Allegra, Jakub Vohryzek, Yonatan Sanz Perl, Nicholas V Metcalf, Joseph C Griffis, Maurizio Corbetta, Gordon L Shulman, and Gustavo Deco. Generative whole-brain dynamics models from healthy subjects predict functional alterations in stroke at the level of individual patients. *Brain Communications*, 6(4):fcae237, 07 2024.